arXiv:1512.00336v2 [stat.AP] 24 Mar 2016

Gaussian and Robust Kronecker Product Covariance Estimation: Existence and Uniqueness

I. Soloveychik[a,*], D. Trushin[a]

[a]*Rachel and Selim Benin School of Computer Science and Engineering, The Hebrew University of Jerusalem, Givat Ram, Jerusalem, Israel, 91904*

Abstract

We study the Gaussian and robust covariance estimation, assuming the true covariance matrix to be a Kronecker product of two lower dimensional square matrices. In both settings we define the estimators as solutions to the constrained maximum likelihood programs. In the robust case, we consider Tyler's estimator defined as the maximum likelihood estimator of a certain distribution on a sphere. We develop tight sufficient conditions for the existence and uniqueness of the estimates and show that in the Gaussian scenario with the unknown mean, $p/q + q/p + 2$ samples are almost surely enough to guarantee the existence and uniqueness, where p and q are the dimensions of the Kronecker product factors. In the robust case with the known mean, the corresponding sufficient number of samples is $\max[p/q, q/p] + 1$.

Keywords: Constrained covariance estimation, robust estimation, high-dimensional estimation, Kronecker product structure.

1. Introduction

Covariance estimation is a fundamental problem in multivariate statistical analysis. It arises in diverse applications such as signal processing, where knowledge of the covariance matrix is unavoidable in constructing optimal detectors Kelly (1986), genomics, where it is widely used to measure correlations between gene expression values Hero and Rajaratnam (2012); Schäfer and Strimmer (2005); Xie and Bentler (2003), and functional MRI

*Corresponding author

Email address: ilya.soloveychik@mail.huji.ac.il (I. Soloveychik)

Derado et al. (2010). Most of the modern algorithms analyzing social networks are based on Gaussian Graphical Models Lauritzen (1996), where the independences between the graph nodes are completely determined by the sparsity structure of the inverse covariance matrix Banerjee et al. (2008). In empirical finance, knowledge of the covariance matrix of stock returns is a fundamental question with implications for portfolio selection and for tests of asset pricing models such as the CAPM Bai and Shi (2011); Ledoit and Wolf (2003). Application of structured covariance matrices instead of Bayesian classifiers based on Gaussian mixture densities or kernel densities proved to be very efficient for many pattern recognition tasks, among them speech recognition, machine translation and object recognition in images Dahmen et al. (2000). In geometric functional analysis and computational geometry Adamczak et al. (2010) the exact estimation of covariance matrix is necessary to efficiently compute volume of a body in high dimension. The classical problems of clustering and Discriminant Analysis are entirely based on precise knowledge of covariance matrices of the involved populations Friedman (1989), etc.

In many modern applications, data sets are very large with both large number of samples n and large dimension p, often with $p \gg n$, leading to the amount of unknown parameters greatly exceeding the number of observations. This high-dimensional regime naturally calls for exploiting or assuming additional structural properties of the data to reduce the number of estimated degrees of freedom. Usually, the specific structures are chosen to be linear or affine. The most popular examples include such models as Toeplitz Abramovich et al. (2007); Asif and Moura (2005); Bickel and Levina (2008); Fuhrmann (1991); Kavcic and Moura (2000); Roberts and Ephraim (2000); Snyder et al. (1989); Soloveychik and Wiesel (2014); Sun et al. (2015); Wiesel et al. (2013), group symmetric Shah and Chandrasekaran (2012); Soloveychik and Wiesel (2016), sparse Banerjee et al. (2008); Ravikumar et al. (2011); Rothman et al. (2008), low rank Fan et al. (2008); Johnstone and Lu (2009); Lounici et al. (2014) and many others. Non-linear structures are also quite common in engineering applications. Among them are the Kronecker product model Dawid (1981); Dutilleul (1999); Tsiligkaridis et al. (2013); Werner et al. (2008), linear and sparse inverse covariance structures such as graphical models Friedman et al. (2008); Yuan and Lin (2007) and others.

In this paper we focus on the Kronecker Product (KP) structure, which has recently become an extremely popular model for a variety of applications, such as MIMO wireless communications Werner and Jansson (2007),

geostatistics Cressie (2015), genomics Yin and Li (2012), multi-task learning Bonilla et al. (2007), face recognition Zhang and Schneider (2010), recommendation systems Allen and Tibshirani (2010), collaborative filtering Yu et al. (2009) and many others. The KP model assumes a $pq \times pq$ covariance matrix $\mathbf{\Theta}_0$ to be a KP of two lower dimensional square matrices, which is denoted by $\mathbf{\Theta}_0 = \mathbf{P} \otimes \mathbf{Q}$, where $\mathbf{P}$ and $\mathbf{Q}$ are $p \times p$ and $q \times q$ dimensional positive definite matrices, respectively. Given $\mathbf{\Theta}_0$, its factors $\mathbf{P}$ and $\mathbf{Q}$ can only be determined up to a positive scalar. This natural ambiguity is usually treated by fixing scaling of one of the factors as we do below.

Consider the Gaussian setting and assume we are given n independent and identically distributed (i.i.d.) pq dimensional real vector measurements $\mathbf{x}_i \sim \mathbf{x},\ i = 1, \dots, n$, where

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{\Theta}). \tag{1}$$

Assume the mean $\boldsymbol{\mu}$ is known, while the covariance $\mathbf{\Theta}$ is to be estimated. If the number of samples is not less than the ambient dimension, $n \geqslant pq$, the Maximum Likelihood Estimator (MLE) of the covariance parameter almost surely exists and coincides with the Sample Covariance Matrix (SCM),

$$\mathbf{S} = \frac{1}{n} \sum_{i=1}^{n} (\mathbf{x}_i - \boldsymbol{\mu})(\mathbf{x}_i - \boldsymbol{\mu})^\top. \tag{2}$$

When the prior knowledge suggests that the true covariance matrix $\mathbf{\Theta}_0$ is of the KP structure, it is usually more convenient to cut $\mathbf{x}$ into q columns of height p each to obtain a so-called matrix normal random variable $\mathbf{X}$ Dawid (1981); Dutilleul (1999); Gupta and Nagar (1999). Following Gupta and Nagar (1999), we denote this law by

$$\mathbf{X} \sim \mathcal{MN}(\mathbf{M}, \mathbf{P} \otimes \mathbf{Q}), \tag{3}$$

where $\mathbf{M}$ is obtained from $\boldsymbol{\mu}$ by the same reshaping procedure. Assume we are given n i.i.d. matrix samples $\mathbf{X}_i \sim \mathbf{X},\ i = 1, \dots, n$ as in (3), and want to estimate the covariance matrix factors $\mathbf{P}$ and $\mathbf{Q}$. Here, the MLE solution is no longer given by an explicit formula as in (2), moreover, the resulting optimization program is non-convex due to the constraint. Luckily, there exists an alternating optimization approach, which is usually adopted Dutilleul (1999); Lu and Zimmerman (2004, 2005); Werner et al. (2008). This algorithm is often referred to as the Flip-Flop (FF) due to the symmetric

updates of the estimates of $\mathbf{P}$ and $\mathbf{Q}$ it produces. Below we show that the obtained constrained program becomes convex under a specific change of metric over the set of positive definite matrices, the so-called geodesic metric Wiesel (2012a,b), naturally explaining the convergence of the FF and significantly helping to further explore the optimization at hand. We refer to this iterative algorithm as the Gaussian FF (GFF) to distinguish it from another FF scheme introduced later.

In many real world applications the underlying multivariate distribution is actually non-Gaussian and robust covariance estimation methods are required. This occurs whenever the distribution of the measurements is heavy-tailed or a small proportion of the samples exhibits outlier behavior Huber (1964); Maronna (1976). Probably the most common extension of the Gaussian family of distributions allowing for treating heavy-tailed populations is the class of elliptically shaped distributions Frahm (2004). Elliptical populations served as the basis for defining a family of the so-called covariance M-estimators Maronna (1976), of which we focus on Tyler's estimator Tyler (1987). Given n samples $\mathbf{x}_i \in \mathbb{R}^{pq}$, $i = 1, \ldots, n$, Tyler's covariance matrix estimator is defined as the solution to the fixed point equation

$$\mathbf{T} = \frac{pq}{n} \sum_{i=1}^{n} \frac{\mathbf{x}_i \mathbf{x}_i^\top}{\mathbf{x}_i^\top \mathbf{T}^{-1} \mathbf{x}_i}. \tag{4}$$

When $\mathbf{x}_i$ are i.i.d. Generalized Elliptically (GE) distributed Frahm (2004), their shape matrix $\mathbf{\Theta}$ is positive definite and $n > pq$, Tyler's estimator exists with probability one and is a consistent estimator of $\mathbf{\Theta}$ up to a positive scaling factor. The GE family includes as particular cases the generalized Gaussian, the compound Gaussian, the elliptical distributions and many others Frahm (2004). Therefore, Tyler's estimator has been successfully used to replace the SCM in many applications sensitive to outliers or heavy-tailed noise, such as anomaly detection in wireless sensor networks Chen et al. (2011), antenna array processing Ollila and Koivunen (2003), and radar detection Abramovich et al. (2007); Bandiera et al. (2010); Ollila et al. (2012); Pascal et al. (2008).

It was recently demonstrated that Tyler's estimator can be viewed as an MLE of a certain spherical distribution Greco and Gini (2013); Soloveychik and Wiesel (2015); Wiesel (2012a). In spite of the fact that the obtained MLE program is not convex, it was later shown to become convex under the geodesic metric change we mentioned above Wiesel (2012a,b). Both these

fundamental discoveries paved a way to the creation of a very useful natural optimization framework characterizing Tyler's estimator, which made possible definition of structured analogs of Tyler's estimator under geodesically convex constraints Soloveychik and Wiesel (2016); Sun et al. (2015); Wiesel and Zhang (2015). The present article extensively uses this new framework to study the existence and uniqueness of the KP constrained Tyler's MLE and the convergence properties of the Robust Flip-Flop (RFF) analog of the GFF algorithm obtained from it. Another very popular in engineering applications example of a linear geodesically convex structure is the so-called group symmetry Shah and Chandrasekaran (2012). Interestingly, a very recent paper Soloveychik and Wiesel (2016) utilized the aforementioned optimization methodology to thoroughly investigate the group symmetric Tyler's estimator (STyler) and its performance benefits. It is important to note that multiple geodesically convex constraints can be efficiently superimposed when the underlying physics suggests such prior knowledge, e.g quite often in practice the KP structure is followed by group symmetries, leading to a further decrease in the number of estimated degrees of freedom.

In both Gaussian and robust cases, one of the central questions in high-dimensional environment is: What is the minimal number of samples guaranteeing the existence and uniqueness of the corresponding covariance MLE? As we have already mentioned, in the unconstrained Gaussian MLE it is known that $n = pq$ samples are enough to guarantee the existence and uniqueness almost surely when the mean is known, and $n = pq + 1$ when the mean is unknown. This number is, of course, enough in the constrained case as well, however, one would expect that this threshold can be reduced due to the decrease in the estimated number of parameters. Different necessary and sufficient conditions on the number of samples in the Gaussian KP scenario were proposed by a large number of works, see Dutilleul (1999); Lu and Zimmerman (2004); Roś et al. (2016); Srivastava et al. (2008); Werner et al. (2008) and references therein. In particular, in Dutilleul (1999) it was claimed that the number of samples needed to guarantee both the existence and uniqueness of the GFF solution in the unknown mean case, equals $\max[p/q, q/p]+1$. Later, Srivastava et al. (2008) showed that, in fact, $\max[p, q] + 1$ matrix valued measurements are required to guarantee the uniqueness, assuming the estimator exists. In Roś et al. (2016) the authors showed by a few simple counterexamples that both results from Dutilleul (1999) and Srivastava et al. (2008) are not correct. Instead, they claimed that "As yet, there do not seem to be existence results for the case $n \in [\max[p/q, q/p] + 1, pq]$,"

therefore, leaving this question open.

Unlike the Gaussian setting, in the robust scenario the mean is usually assumed to be known. To the best of our knowledge, in this case the question of the minimal number of samples needed to ensure existence and uniqueness was not properly addressed thus far, except for the trivial necessary condition $n \geqslant \max[p/q, q/p]$ stemming from the definition of RFF, as shown below.

The main goal of this paper is to present tight thresholds for the necessary and sufficient conditions guaranteeing existence and uniqueness in both Gaussian and robust cases. Namely, we show that in the Gaussian setting, when the mean is not known, if $n < \max[p/q, q/p] + 1$, the estimator, even if it exists, is not unique with probability one and, when $n > p/q + q/p + 1$, the solution exists and is unique almost surely. We also provide a discussion explaining that between these bounds the probabilities of existence and non-existence of a unique solution are positive. In the robust case with the mean known, the threshold is $\max[p/q, q/p]$. More specifically, if n is less than this number, no unique solution exists, while if n is greater than this value, the estimator almost surely exists and is unique.

The rest of the text is organized as follows. After we introduce notations, we define the Gaussian setting and formulate the problem. Then we discuss the state of the art results concerning the necessary and sufficient conditions for the existence and uniqueness. Section 4 presents our main result, shows the main idea of the proof and demonstrates it on a simple two dimensional example. Then we treat the robust scenario and provide our main contribution there. Finally, Sections 6 and 7 contain all the proofs for the Gaussian and robust cases, correspondingly.

1.1. Notations

We deal with real Euclidean spaces denoted by $\mathbb{R}^p$, whose elements are columns written as lower case bolds $\mathbf{x}$. The standard scalar products over such spaces are denoted by $(\cdot, \cdot)$ and the corresponding norms by $\|\cdot\|$. By $\mathcal{M}_{p\times q}$ we denote the Euclidean space of real $p \times q$ matrices, written as upper case bolds $\mathbf{X}$. $\mathcal{S}(\mathbb{R}^p)$ stands for the linear space of symmetric $p \times p$ matrices and $\mathcal{P}(\mathbb{R}^p) \subset \mathcal{S}(\mathbb{R}^p)$ - for the open cone of positive definite matrices inside it. Note that $\mathcal{S}(\mathbb{R}^p)$ inherits from $\mathcal{M}_{p\times p}$ the natural structure of a Euclidean space with the Frobenius norm. $\mathbf{I}$ denotes the identity matrix of an appropriate dimension. For two spaces $\mathbb{R}^p$ and $\mathbb{R}^q$, $\mathbb{R}^p \otimes \mathbb{R}^q$ denotes their tensor product space and for two matrices $\mathbf{P} \in \mathcal{P}(\mathbb{R}^p)$, $\mathbf{Q} \in \mathcal{P}(\mathbb{R}^q)$, $\mathbf{P} \otimes \mathbf{Q}$ denotes their Kronecker product. The spectral norm of a matrix $\mathbf{P} \in \mathcal{P}(\mathbb{R}^p)$ is

denoted by $\|\mathbf{P}\|_2$ and its determinant by $|\mathbf{P}|$. Given a set X, its boundary is denoted by ∂X. For two sets X and Y, $X \times Y$ denotes their Cartesian (direct) product. Given a subset X of a linear space, $\langle X \rangle$ denotes its span and $|X|$ - its cardinality. We use standard abbreviation a.s. to denote the almost sure convergence when the measure can be inferred from the context. The symbol $\sim$ replaces saying "is distributed identically to".

2. Gaussian Setting and Problem Formulation

Assume we are given n i.i.d. Gaussian matrix samples

$$\mathbf{X}_1, \ldots, \mathbf{X}_n \sim \mathbf{X}, \qquad \mathbf{X} \sim \mathcal{MN}(\mathbf{M}, \mathbf{P} \otimes \mathbf{Q}), \tag{5}$$

where $\mathbf{X}_i \in \mathcal{M}_{p\times q}$, $\mathbf{P} \in \mathcal{P}(\mathbb{R}^p)$ and $\mathbf{Q} \in \mathcal{P}(\mathbb{R}^q)$. Denote $X = \{\mathbf{X}_1, \ldots, \mathbf{X}_n\}$, then, up to an additive constant and scaling, the negative log-likelihood reads as

$$\widetilde{f}_{\mathcal{N}}(\mathbf{M}, \mathbf{P} \otimes \mathbf{Q}; X) = \frac{1}{n} \sum_{i=1}^{n} \operatorname{Tr}\left(\mathbf{P}^{-1}(\mathbf{X}_i - \mathbf{M})\mathbf{Q}^{-1}(\mathbf{X}_i - \mathbf{M})^\top\right) + \ln|\mathbf{P} \otimes \mathbf{Q}|, \tag{6}$$

and is defined over the set $\mathcal{M}_{p\times q} \times \mathcal{M}_{\mathcal{N}}$, with

$$\mathcal{M}_{\mathcal{N}} = \{\mathbf{P} \otimes \mathbf{Q} \mid \mathbf{P} \in \mathcal{P}(\mathbb{R}^p),\ \mathbf{Q} \in \mathcal{P}(\mathbb{R}^q)\} \subset \mathcal{P}(\mathcal{M}_{p\times q}). \tag{7}$$

Here, the matrix $\mathbf{P} \otimes \mathbf{Q}$ is identified with the positive operator $\mathbf{P} \otimes \mathbf{Q} : \mathcal{M}_{p\times q} \to \mathcal{M}_{p\times q}$ acting by the rule $\mathbf{X} \to \mathbf{PXQ}$. The scalar product on $\mathcal{M}_{p\times q}$ is given by $(\mathbf{A}, \mathbf{B}) = \operatorname{Tr}\left(\mathbf{AB}^\top\right)$. Note that $\mathcal{M}_{\mathcal{N}}$ can be identified with the set

$$\mathcal{M}_{\mathcal{N}} \cong \{(\mathbf{P}, \mathbf{Q}) \mid \|\mathbf{P}\|_2 = 1\} \subset \mathcal{P}(\mathbb{R}^p) \times \mathcal{P}(\mathbb{R}^q), \tag{8}$$

where the specific normalization can be chosen arbitrarily.

Remark 1. *Below we assume the following notational convention: when the set $\mathcal{M}_{\mathcal{N}}$ is viewed as a subspace of $\mathcal{P}(\mathcal{M}_{p\times q})$ as in (7), the arguments of the negative log-likelihood are written as $\widetilde{f}_{\mathcal{N}}(\mathbf{M}, \mathbf{P} \otimes \mathbf{Q}; X)$, however, when $\mathcal{M}_{\mathcal{N}}$ is identified with a subset of $\mathcal{P}(\mathbb{R}^p) \times \mathcal{P}(\mathbb{R}^q)$ defined by (8), we write the arguments as $\widetilde{f}_{\mathcal{N}}(\mathbf{M}, \mathbf{P}, \mathbf{Q}; X)$, with*

$$\widetilde{f}_{\mathcal{N}}(\mathbf{M}, \mathbf{P}, \mathbf{Q}; X) = \widetilde{f}_{\mathcal{N}}(\mathbf{M} \otimes \mathbf{P}, \mathbf{Q}; X). \tag{9}$$

Below we use the same rule for other similar functions and the specific representation of the underlying set can be inferred from the way arguments are written.

Identification (8) allows us to consider elements $\mathbf{P} \otimes \mathbf{Q} \in \mathcal{M}_{\mathcal{N}}$ (under proper normalization $\mathbf{P} \otimes \mathbf{Q} = (\mathbf{P}/\|\mathbf{P}\|_2) \otimes (\|\mathbf{P}\|_2 \mathbf{Q})$ if needed) as pairs $(\mathbf{P}, \mathbf{Q})$. In addition, it endows $\mathcal{M}_{\mathcal{N}}$ with a smooth manifold structure making $\widetilde{f}_{\mathcal{N}}$ a smooth function over it. The covariance MLE under the KP constraint can now be written as a solution to the following program

$$\min_{\substack{\mathbf{M} \in \mathcal{M}_{p \times q}, \\ (\mathbf{P}, \mathbf{Q}) \in \mathcal{M}_{\mathcal{N}}}} \widetilde{f}_{\mathcal{N}}(\mathbf{M}, \mathbf{P}, \mathbf{Q}; X). \tag{10}$$

As in the unconstrained Gaussian case, this program decouples into minimization w.r.t. (with respect to) the unknown mean $\mathbf{M}$, yielding

$$\widehat{\mathbf{M}} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{X}_i, \tag{11}$$

and minimization w.r.t. to $\mathbf{P}$ and $\mathbf{Q}$. Note that $\ln |\mathbf{P} \otimes \mathbf{Q}| = q \ln |\mathbf{P}| + p \ln |\mathbf{Q}|$, and denote $\mathbf{Y}_i = \mathbf{X}_i - \widehat{\mathbf{M}}$, then the first-order optimal conditions for $\mathbf{P}$ and $\mathbf{Q}$ read as

$$\begin{cases} \mathbf{P} = \dfrac{1}{qn} \displaystyle\sum_{i=1}^{n} \mathbf{Y}_i \mathbf{Q}^{-1} \mathbf{Y}_i^{\top}, \\ \mathbf{Q} = \dfrac{1}{pn} \displaystyle\sum_{i=1}^{n} \mathbf{Y}_i^{\top} \mathbf{P}^{-1} \mathbf{Y}_i. \end{cases} \tag{12}$$

There does not exist a closed form analytic solution to (12), therefore, it is usually solved numerically via the so-called Flip-Flop (FF) iterative scheme Dutilleul (1999), which we call the Gaussian FF (GFF). The GFF algorithm works as follows. Starting from an initial guess $(\mathbf{P}_0, \mathbf{Q}_0) \in \mathcal{M}_{\mathcal{N}}$ for $(\mathbf{P}, \mathbf{Q})$, we plug it into the right-hand side of (12) and get a new pair $(\mathbf{P}_1, \mathbf{Q}_1)$. After we normalize this pair to make it formally belong to $\mathcal{M}_{\mathcal{N}}$, we apply the procedure to $(\mathbf{P}_1, \mathbf{Q}_1)$ instead $(\mathbf{P}_0, \mathbf{Q}_0)$ and so on as shown in diagram (17),

$$\begin{cases} \widetilde{\mathbf{P}}_{j+1} = \dfrac{1}{qn} \displaystyle\sum_{i=1}^{n} \mathbf{Y}_i \mathbf{Q}_j^{-1} \mathbf{Y}_i^{\top}, \\ \widetilde{\mathbf{Q}}_{j+1} = \dfrac{1}{pn} \displaystyle\sum_{i=1}^{n} \mathbf{Y}_i^{\top} \mathbf{P}_j^{-1} \mathbf{Y}_i, \\ \mathbf{P}_{j+1} = \dfrac{1}{\|\widetilde{\mathbf{P}}_{j+1}\|_2} \widetilde{\mathbf{P}}_{j+1}, \\ \mathbf{Q}_{j+1} = \|\widetilde{\mathbf{P}}_{j+1}\|_2 \widetilde{\mathbf{Q}}_{j+1}. \end{cases} \tag{13}$$

The consecutive pairs $(\mathbf{P}_j, \mathbf{Q}_j)$ serve as approximations to the true solution, therefore, the convergence of this sequence to the minimum of $\widetilde{f}_{\mathcal{N}}(\widehat{\mathbf{M}}, \mathbf{P}, \mathbf{Q}; X)$ on $\mathcal{M}_{\mathcal{N}}$ (if it exists) constitutes one of the central topics of our paper. We start form listing the existing results on the questions at hand.

3. Existence, Uniqueness and Convergence: State of the Art

Having derived the G-CARMEL (Gaussian KRonecker product MLE) solution and obtained an iterative scheme for its calculation, our next goal is to determine the necessary and sufficient conditions for its existence and uniqueness and for the convergence of the GFF procedure. The only parameter under our control is the required number of i.i.d. normal samples, therefore, below we focus on the question: How many measurements one needs to guarantee existence, uniqueness and convergence almost surely?

- **Existence**. We start from the sufficient conditions. It was claimed in Dutilleul (1999) that $\max[p/q, q/p] + 1$ samples are needed for the existence and uniqueness of the MLE solution in the Gaussian case. However, it was later shown by a counterexample Roś et al. (2016) that the uniqueness does not follow from this condition. In addition, the authors of Roś et al. (2016) write that "Moreover, it is not known whether it [this condition] guaranties existence, because it is not sufficient to show that all updates of the FF algorithm have full rank as is done in Dutilleul (1999). It could still happen that the sequence of updates converges (after infinitely many steps) to a Kronecker product that does not have a full rank with the likelihood converging to its supremum." It is also claimed in Roś et al. (2016) that no less than $pq+1$ samples are required to ensure the existence a.s. This number of measurements coincides with the one needed in the unconstrained case and does not explore the KP structure. Finally, the authors of Roś et al. (2016) conclude that nothing can be said regarding the existence, if the number of samples lies inside the interval $n \in [\max[p/q, q/p + 1, pq]$.

  The necessary conditions were also treated in Dutilleul (1999), where the author claims that if the estimator exists, then $n \geqslant \max[p/q, q/p]+1$. This is clearly true, since if the number of samples is less than this threshold, at least one of the right-hand sides in (12) is rank deficient and cannot be invertible.

- **Uniqueness**. As summarized in Roś et al. (2016), the author of Dutilleul (1999) claims that the G-CARMEL is unique whenever $n \geqslant \max[p/q, q/p]+$

1. Later, the authors of Srivastava et al. (2008) stated that indeed $n \geqslant \max[p, q] + 1$ is needed to ensure the uniqueness. Here again, Roś et al. (2016) succeeded to find counterexamples showing that both these bounds do not guarantee uniqueness. Moreover, the paper Roś et al. (2016) describes the exact parts of the proofs which seem to contain mistakes, however, the correct lower bounds on the number of required samples are not provided. In fact, to the best of our knowledge, tight sufficient conditions for the uniqueness have not been reported so far.

- **Convergence of the Gaussian Flip-Flop Algorithm**. The last question regarding the G-CARMEL on which we focus, is the convergence of the GFF iterative scheme. In Dutilleul (1999) the author establishes the convergence of the GFF technique empirically. He claims that if the limiting points of the sequences $\{\widehat{\mathbf{Q}}_j\}$ and $\{\widehat{\mathbf{P}}_j\}$ (in his notations) do not depend on the initial point, and an additional condition on the second derivatives of the objective is satisfied at the limiting points, then these limits provide the G-CARMEL solution. If such limiting points are not uniquely determined, but rather depend on the initial guesses, they must provide local extrema of the likelihood function. Unfortunately, this empirical approach can hardly be applied in practice and does not provide a strict criterion for the convergence of the GFF.

  The authors of Lu and Zimmerman (2004, 2005) claim that when the number of samples is $n \geqslant pq + 1$, the GFF is guarantied to converge, however they doubt if it really converges to the MLE, since the "parameter space of (p, q)-separable covariance matrices is not convex". They emphasize that for some values of n the algorithm can converge to many different estimates, depending on the starting value. Finally, they conjecture that for n large enough "the limit point of the GFF can safely be regarded as the unique MLE" without proving this statement.

  In Werner et al. (2008) theoretical asymptotic properties of the GFF algorithm are considered and the algorithm's performance for small sample sizes is investigated with a simulation study.

The main contributions of the Gaussian part of our paper consist in

- proving tight sufficient and necessary conditions for the a.s. existence and uniqueness of the G-CARMEL estimate,

- showing that the sufficient conditions also imply convergence of the GFF iterations to the unique solution starting from any initial guess.

4. Main Results and Arguments

In this section we state our main result in the normal case, give the intuition behind the proof argument and demonstrate our technique on a simple example in a low dimension.

4.1. The Main Statement

Theorem 1. *Assume* $X = \{\mathbf{X}_1, \dots \mathbf{X}_n\}$ *are independently sampled from a continuous distribution over* $\mathcal{M}_{p\times q}$ *and consider the problem of minimizing* $\widetilde{f}_{\mathcal{N}}(\mathbf{M}, \mathbf{P}, \mathbf{Q}; X)$ *over* $\mathcal{M}_{p\times q} \times \mathcal{M}_{\mathcal{N}}$*, then*

1. *if* $n < \max[p/q, q/p] + 1$*, there is no unique minimum,*
2. *if* $n > p/q + q/p + 1$*, there is a unique minimum a.s.,*
3. *if* $n > p/q + q/p + 1$*, the GFF converges starting from any point of* $\mathcal{M}_{\mathcal{N}}$ *to this unique minimum a.s.*

Proof. This is a direct corollary of Theorem 4 from Section 6.7. □

Remark 2. *The statement of the theorem is valid for any continuous distribution and is not limited to the Gaussian ones. Indeed, the claim does not assume any specific statistical model and does not provide statistical guaranties (e.g. consistency), but rather treats the questions of the existence and uniqueness of the minimum.*

Remark 3. *Note the gap between items 1) and 2) containing one (when* $p \neq q$*) or two (when* $p = q$*) integer points which cannot be eliminated. We discuss this phenomenon below in more detail.*

4.2. Sketch of the The Proof

In this section we discuss the main building blocks of the proof of Theorem 1 leaving the technical details to Section 6.

- **Reduction to the Centered Case**. Let

$$\widetilde{g}_{\mathcal{N}}(\mathbf{P}\otimes\mathbf{Q};X)=\widetilde{f}_{\mathcal{N}}(0,\mathbf{P}\otimes\mathbf{Q};X), \tag{14}$$

then minimization of $g_{\mathcal{N}}$ over $\mathcal{M}_{\mathcal{N}}$ does not require optimization w.r.t. the mean parameter. The general case with the unknown expectation can be reduced to it through the following observation. Given a family $X=\{\mathbf{X}_1,\ldots,\mathbf{X}_n\}\subset\mathcal{M}_{p\times q}$ of n random $p\times q$ matrices, there always exists another family $Y=\{\mathbf{Y}_1,\ldots,\mathbf{Y}_{n-1}\}\subset\mathcal{M}_{p\times q}$ of $n-1$ random matrices, such that

$$\widetilde{g}_{\mathcal{N}}(\mathbf{P}\otimes\mathbf{Q};Y)=\widetilde{f}_{\mathcal{N}}(\widehat{\mathbf{M}},\mathbf{P}\otimes\mathbf{Q};X). \tag{15}$$

Lemma 11 from Section 6 shows why this is true and justifies our transition to the zero mean case. In the remainder of this section we treat the zero mean setting.

- **Necessary Conditions**. Since we require the solution to be composed of invertible matrices $\mathbf{P}$ and $\mathbf{Q}$, (12) must hold at the extremum point. Note that its right-hand side is not invertible for $n<\max[p/q,q/p]$, therefore, returning to the non-centered case and compensating for this by adding one sample, yields item 1) of Theorem 1.

- **Sufficient Conditions**. To derive the sufficient conditions, in Section 6 we change the parametrization of $\widetilde{g}_{\mathcal{N}}$ by

$$g_{\mathcal{N}}(\mathbf{P}\otimes\mathbf{Q};X)=\widetilde{g}_{\mathcal{N}}(\mathbf{P}^{-1}\otimes\mathbf{Q}^{-1};X)=\frac{1}{n}\sum_{i=1}^{n}(\mathbf{P}\mathbf{X}_i\mathbf{Q},\mathbf{X}_i)-\ln|\mathbf{P}\otimes\mathbf{Q}|\,, \tag{16}$$

and introduce a specific metric over $\mathcal{P}(\mathbb{R}^p\times\mathbb{R}^q)$, w.r.t. which the set $\mathcal{M}_{\mathcal{N}}$ and the function $g_{\mathcal{N}}(\mathbf{P}\otimes\mathbf{Q};X)$ are convex. The desired solution exists and is unique if and only if $g_{\mathcal{N}}$ continuously tends to $+\infty$ on the boundary as shown in Lemma 4, in which case it is also strictly convex. Theorem 3 then demonstrates that this happens a.s. w.r.t. the distribution of X when $n>p/q+q/p$. In the next section we demonstrate the reasoning behind these claims by exploiting the $p=q=2$ case in more detail.

- **Convergence of the GFF**. Suppose we are given a pair of matrices $(\mathbf{P}_0, \mathbf{Q}_0) \in \mathcal{M}_{\mathcal{N}}$ and use (13) to generate the sequence

$$\begin{array}{ccccccc} \mathbf{P}_0 & & \mathbf{P}_1 & & \mathbf{P}_2 & & \dots \\ & \searrow\!\!\nearrow & & \searrow\!\!\nearrow & & \searrow\!\!\nearrow & \\ \mathbf{Q}_0 & & \mathbf{Q}_1 & & \mathbf{Q}_2 & & \dots \end{array} \tag{17}$$

  Here, the successive iterates $\mathbf{P}_j, \mathbf{Q}_{j+1}, \mathbf{P}_{j+2}, \dots$ are obtained by minimizing $g_{\mathcal{N}}$ w.r.t. $\mathbf{Q}$ when $\mathbf{P}_j$ is fixed and similarly by minimizing w.r.t. $\mathbf{P}$ when $\mathbf{Q}_j$ is fixed, etc. As we have mentioned in the previous paragraph, $g_{\mathcal{N}}$ is a.s. strictly convex and tends to $+\infty$ on the boundary when $n > p/q + q/p$. This guarantees a decrease of the target function on each iteration and the convergence of the sequence $(\mathbf{P}_j, \mathbf{Q}_{j+1})$ to the unique minimum, hence, $(\mathbf{P}_j, \mathbf{Q}_j)$ converges as well.

Let us now illustrate the main arguments by a simple low dimensional example.

4.3. $p = q = 2$ Case Study

Assume $\mathbb{R}^p = \mathbb{R}^q = \mathbb{R}^2$, $X = \{\mathbf{X}_1, \dots, \mathbf{X}_n\} \subset \mathcal{M}_{2\times 2}$ consists of n matrices and we deal with the case of zero mean. We are going to show a bit more than we have announced in the previous section, namely, we will prove that

1. If $n = 1$, the set of minima is non empty and forms a submanifold of dimension 3 with probability one. In particular, a minimum exists but is not unique.

2. If $n = 2$, there exists a polynomial $D(\mathbf{X}_1, \mathbf{X}_2)$ such that

   - if $D(\mathbf{X}_1, \mathbf{X}_2) \geqslant 0$, there is no unique minimum of $g_{\mathcal{N}}$ over $\mathcal{M}_{\mathcal{N}}$,
   - if $D(\mathbf{X}_1, \mathbf{X}_2) < 0$, there is a unique minimum of $g_{\mathcal{N}}$ over $\mathcal{M}_{\mathcal{N}}$,

   both happening with positive probabilities.

3. If $n > 2$, there is a unique minimum of $g_{\mathcal{N}}$ over $\mathcal{M}_{\mathcal{N}}$.

As explained in Section 6.2, the set $\mathcal{M}_{\mathcal{N}} \subset \mathcal{P}(\mathbb{R}^2 \times \mathbb{R}^2)$ is convex w.r.t. to a specific metric change. In addition, Lemma 4 demonstrates that the solution to the optimization at hand exists and is unique if and only if $g_{\mathcal{N}}$ continuously tends to $+\infty$ on the boundary, which we use below.

1) $n = 1$ **case.** Here, both equations in (12) (after replacing $\mathbf{P}$ and $\mathbf{Q}$ by their inverses) become identical to

$$\mathbf{P}^{-1} = \frac{1}{2}\mathbf{X}_1\mathbf{Q}\mathbf{X}_1^\top. \tag{18}$$

This equation defines a submanifold $\mathcal{M}_m \subset \mathcal{M}_\mathcal{N}$ isomorphic to $\mathcal{P}(\mathbb{R}^2)$ containing $\mathbf{Q}$-s, and, therefore, having dimension 3. A straightforward computation shows that the value of $g_\mathcal{N}$ is constant on $\mathcal{M}_m$. Since $g_\mathcal{N}$ is convex, all points of $\mathcal{M}_m$ are minima.

2) $n = 2$ **case.** It turns out that the critical question defining the behavior of the solution here is whether there exists a vector $\mathbf{t} \in \mathbb{R}^2$ such that $\mathbf{X}_1\mathbf{t}$ and $\mathbf{X}_2\mathbf{t}$ are parallel. If the answer is negative, the minimum exists and is unique, otherwise, if it exists, it is not unique. As Lemma 4 item 4) shows below, such vector $\mathbf{t}$ does not exist if and only if $g_\mathcal{N}$ tends to $+\infty$ on the boundary of $\mathcal{M}_\mathcal{N}$. Next we explain the reasoning in more detail and explicitly construct such $\mathbf{t}$.

Consider a sequence $\mathcal{M}_\mathcal{N} \ni \{(\mathbf{P}_j, \mathbf{Q}_j)\} \to \partial\mathcal{M}_\mathcal{N}$, meaning that either $\mathbf{P}_j$ tend to a singular matrix or the norms of $\mathbf{Q}_j$ are unbounded (or both). Below we suppress the j indexing of the sequences to simply notations. In other words we distinguish between two cases: a) either $\|\mathbf{P}\|_2$ and $\|\mathbf{Q}\|_2$ are bounded or b) we may assume that in some appropriately chosen bases $\{\mathbf{s}_1, \mathbf{s}_2\}$ and $\{\mathbf{t}_1, \mathbf{t}_2\}$ for $\mathbf{P}$ and $\mathbf{Q}$, respectively, we have

$$\mathbf{P} = \begin{pmatrix} \alpha & 0 \\ 0 & 1 \end{pmatrix} \quad \mathbf{Q} = \begin{pmatrix} \mu & 0 \\ 0 & \eta \end{pmatrix}, \tag{19}$$

where $\alpha \leqslant 1$ and we assume (after swapping the eigenbasis of $\mathbf{Q}$, if necessary) that $\mu \to +\infty$ not slower than η (if η is bounded this is vacuously true).

In the first case, when the spectral norms are bounded, the trace term of $g_\mathcal{N}$ is bounded. Since in this scenario at least one of the matrices must tend to a singular one, $\ln|\mathbf{P} \otimes \mathbf{Q}| \to -\infty$, implying $g_\mathcal{N} \to +\infty$. In the second case we have sequence (19) and note that the logarithmic term of $g_\mathcal{N}$ has a summand tending to $-\infty$ with the rate not greater than $\ln \mu$. Assume $\mathbf{X}_1\mathbf{t}_1$ and $\mathbf{X}_2\mathbf{t}_1$ are not parallel for all $\mathbf{t}_1$, then at least one of $\mathbf{X}_i$-s has a non zero $(2,1)$ element. Suppose this is $\mathbf{X}_1 = \left(\begin{smallmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{smallmatrix}\right)$ with $x_{21} \neq 0$, then the scalar products part of $g_\mathcal{N}$ is not less than $\frac{1}{2}|x_{21}|^2\mu$. Hence, this part of $g_\mathcal{N}$ tends to $+\infty$ faster than the negative part and totally $g_\mathcal{N} \to +\infty$ on the sequence at hand.

Now suppose that there does exist a vector $\mathbf{t} \in \mathbb{R}^2$ such that $\mathbf{X}_1\mathbf{t}$ and $\mathbf{X}_2\mathbf{t}$ are collinear. Normalize $\mathbf{t}$ and form an orthonormal basis $\{\mathbf{t}, \mathbf{t}'\}$ in the space of $\mathbf{Q}$. After this, normalize $\mathbf{X}_1\mathbf{t}$, which we denote by $\mathbf{s}$, and complete it to an orthonormal basis $\{\mathbf{s}, \mathbf{s}'\}$ in the space of $\mathbf{P}$. In these bases each $\mathbf{X}_i$ reads as (here we omit index i in matrix elements for simplicity)

$$\mathbf{X}_i = \begin{pmatrix} x_{11} & x_{12} \\ 0 & x_{22} \end{pmatrix}. \tag{20}$$

Now define a new sequence in the chosen bases

$$\mathbf{P} = \begin{pmatrix} \frac{1}{\mu} & 0 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{Q} = \begin{pmatrix} \mu & 0 \\ 0 & 1 \end{pmatrix}, \tag{21}$$

with $\mu \to +\infty$. Then $\mathbf{P} \otimes \mathbf{Q}$ tends to the boundary of $\mathcal{M}_{\mathcal{N}}$,

$$\ln|\mathbf{P} \otimes \mathbf{Q}| = 2\ln|\mathbf{P}| + 2\ln|\mathbf{Q}| = 2\ln|\mathbf{P}\mathbf{Q}| = 0, \tag{22}$$

and for each $\mathbf{X}_i$,

$$(\mathbf{P}\mathbf{X}_i\mathbf{Q}, \mathbf{X}_i) = x_{12}^2 + x_{21}^2 + x_{22}^2\frac{1}{\mu}. \tag{23}$$

Hence, $g_{\mathcal{N}}$ is bounded on the sequence $\{(\mathbf{P}, \mathbf{Q})\}$ and we are done with this case.

Next we derive a condition on $\mathbf{X}_1$ and $\mathbf{X}_2$ telling whether such a mutual $\mathbf{t}$ exists, which will suggest us the probability of such event. Let in the original bases $\mathbf{X}_1$, $\mathbf{X}_2$ and $\mathbf{t}$ read as

$$\mathbf{X}_1 = \begin{pmatrix} x & y \\ u & v \end{pmatrix}, \quad \mathbf{X}_2 = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \quad \mathbf{t} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}. \tag{24}$$

We look for all triples $(\mathbf{X}_1, \mathbf{X}_2, \mathbf{t})$ such that $\mathbf{X}_1\mathbf{t}$ and $\mathbf{X}_2\mathbf{t}$ are collinear, which is equivalent to

$$\begin{vmatrix} \alpha x + \beta y & \alpha a + \beta b \\ \alpha u + \beta v & \alpha c + \beta d \end{vmatrix} = 0. \tag{25}$$

The latter can be written as

$$\begin{vmatrix} x & a \\ u & c \end{vmatrix}\alpha^2 + \left(\begin{vmatrix} x & b \\ u & d \end{vmatrix} + \begin{vmatrix} y & a \\ v & c \end{vmatrix}\right)\alpha\beta + \begin{vmatrix} y & b \\ v & d \end{vmatrix}\beta^2 = 0. \tag{26}$$

Calculate the discriminant of this quadric

$$D(\mathbf{X}_1, \mathbf{X}_2) = \left(\begin{vmatrix} x & b \\ u & d \end{vmatrix} + \begin{vmatrix} y & a \\ v & c \end{vmatrix}\right)^2 - 4\begin{vmatrix} x & a \\ u & c \end{vmatrix}\begin{vmatrix} y & b \\ v & d \end{vmatrix}. \tag{27}$$

Note that $D(\mathbf{X}_1, \mathbf{X}_2)$ is a non-zero polynomial and there is $\mathbf{t} \in \mathbb{R}^2$ with the required properties if and only if $D \geqslant 0$. If the density of the distribution of $(\mathbf{X}_1, \mathbf{X}_2)$ is a.s. non-zero, then clearly $D \geqslant 0$ and $D < 0$ both hold with non-zero probabilities.

3) $n > 2$ **case.** Here a similar computation shows that we a.s. cannot find a vector $\mathbf{t}$ such that $\mathbf{X}_i\mathbf{t}$ are collinear, therefore the above arguments imply the existence and uniqueness of the minimum.

4.4. Remarks

To summarize, the answer to the existence and uniqueness question can be completely described in terms of the following indicator variable:

$$\zeta(X) = \operatorname*{ess\,inf}_{\mathbf{u} \in \mathbb{R}^2 \setminus 0} \dim \left(\sum_{\mathbf{X}_i \in X} \mathbf{X}_i \langle \mathbf{u} \rangle \right). \tag{28}$$

When $n = 1$, $\zeta(X) = 1$ a.s., in the case $n > 2$, $\zeta(X) = 2$ a.s., and these two situations correspond to the uniqueness and non uniqueness. When $n = 2$, we have

$$\zeta(X) = \begin{cases} 1, & \text{if } D \geqslant 0, \\ 2, & \text{if } D < 0, \end{cases} \tag{29}$$

where both events happen with non-zero probabilities, i.e. $\zeta(X)$ is not a constant a.s. This intuitively explains the one sample gap between the necessary and sufficient conditions in Theorem 1.

In arbitrary dimension the ideas described above generalize as following. In order to guarantee the desired asymptotic behavior of the target function, our aim would be to avoid the following situation: there is a random subspace $U \subseteq \mathbb{R}^q$ such that the dimension of $\sum_{i=1}^n \mathbf{X}_i U$ is less than $\min[n \dim U, p]$ with non-zero probability. As the proof of Theorem 1 shows, when the number of samples satisfies the required condition, such event will a.s. not happen.

Let us focus on the gap between the necessary and sufficient conditions appearing in the statement of Theorem, 1

$$\mathcal{I} = \left[\max\left(\frac{p}{q}, \frac{q}{p}\right) + 1, \quad \frac{p}{q} + \frac{q}{p} + 1 \right]. \tag{30}$$

This interval contains 2 points in case $p = q$ and only 1 point otherwise. The same argument as in the example above (note that in the example we considered the known mean case, therefore the $n = 1$ case considered

there would correspond to $n = 2$ here) shows that when $p = q$ and $n = \max[p/q, q/p] + 1 = 2 \in \mathcal{I}$, there are multiple minima. Therefore, there is only one untreated integer point $p/q+q/p+1$ left inside the interval $\mathcal{I}$. However, we do not investigate deeply the behavior of this remaining value due to the following reason. As the two dimensional example above suggests (this corresponds to the case $n = 2$ case in the example), in this case uniqueness and non-uniqueness happen with non-zero probabilities, making the analysis hard. Since there is only one untreated point left and the treatment involves quite non-trivial calculations, the game does not worth the candle. We believe that in general dimensions this missing point of the interval exhibits the same behavior, and both events "existence and uniqueness" and "existence and non-uniqueness" happen with non-zero probabilities.

5. Robust Kronecker Product Covariance Estimation

5.1. Tyler's Estimator

As we have already explained in the Introduction, when robust covariance estimation is considered, the most popular tool used by practitioners is the so-called family of covariance M-estimators introduced by R.A. Maronna Maronna (1976). We focus on a distribution-free member of this class introduced by D.E. Tyler and named after him Tyler (1987). Tyler's covariance estimator, given by formula (4), can be equivalently defined as a covariance parameter MLE of a certain spherical distribution Greco and Gini (2013); Soloveychik and Wiesel (2015) as follows.

Definition 1. *Assume* $\mathbf{\Theta}_0 \in \mathcal{P}(\mathbb{R}^p)$*, then*

$$p(\mathbf{x}) = \frac{\Gamma(p/2)}{2\sqrt{\pi}^p} \frac{1}{\sqrt{|\mathbf{\Theta}_0|}(\mathbf{x}^H \mathbf{\Theta}_0^{-1} \mathbf{x})^{p/2}} \tag{31}$$

is a probability density function of a vector $\mathbf{x} \in \mathbb{R}^p$ *lying on a unit sphere. This distribution is usually referred to as the Real Angular Central Elliptical (RACE) distribution Greco and Gini (2013), and we denote it as* $\mathbf{x} \sim \mathcal{U}(\mathbf{\Theta}_0)$*. The matrix* $\mathbf{\Theta}_0$ *is referred to as a shape matrix of the distribution.*

RACE distribution is closely related to the class of Generalized Elliptical (GE) populations, which includes Gaussian, compound Gaussian, elliptical, skew-elliptical, RACE and other distributions Frahm and Jaekel (2007). An

important property of the GE family is that the shape matrix of a population does not change when the vector is divided by its Euclidean norm Frahm (2004); Frahm and Jaekel (2007). After normalization any GE vector becomes RACE distributed. This allows us to treat all these distributions together using Tyler's estimator, which is the MLE of the shape matrix parameter in RACE populations and is unbiased when a specific scaling is fixed Greco and Gini (2013); Soloveychik and Wiesel (2015).

5.2. Robust Setting and Problem Formulation

In order to proceed to the KP structured robust covariance estimation, we introduce the following setting. Assume we are given n i.i.d. centered real $p \times q$ matrix measurements $X = \{\mathbf{X}_1, \ldots, \mathbf{X}_n\}$ and our goal is to determine what is the minimal number of samples n needed to ensure the existence and uniqueness of Tyler's estimator under the KP constraint. We use the MLE formulation of Tyler's estimator and consider the corresponding optimization program. Specifically, we search for positive definite $\mathbf{P}$ and $\mathbf{Q}$ minimizing the target

$$\widetilde{f}_{\mathcal{E}}(\mathbf{P} \otimes \mathbf{Q}; X) = \frac{1}{pq} \ln |\mathbf{P} \otimes \mathbf{Q}| + \frac{1}{n} \sum_{i=1}^{n} \ln \left(\mathrm{Tr} \left(\mathbf{P}^{-1} \mathbf{X}_i \mathbf{Q}^{-1} \mathbf{X}_i^{\top} \right) \right), \tag{32}$$

which is a robust version of the G-CARMEL estimator and is named R-CARMEL.

The target $\widetilde{f}_{\mathcal{E}}$ is naturally defined over $\mathcal{M}_{\mathcal{N}}$ introduced in (7) and Remark 1 applies here as well, therefore, we use the same notational convention. In addition, $\widetilde{f}_{\mathcal{E}}$ is scale invariant $\widetilde{f}_{\mathcal{E}}(\lambda \mathbf{P} \otimes \mathbf{Q}; X) = \widetilde{f}_{\mathcal{E}}(\mathbf{P} \otimes \mathbf{Q}; X)$, hence, we rather consider $\widetilde{f}_{\mathcal{E}}$ over

$$\mathcal{M}_{\mathcal{E}} = \mathcal{M}_{\mathcal{N}} / \{\mathbf{P} \otimes \mathbf{Q} \sim \lambda \mathbf{P} \otimes \mathbf{Q},\ \lambda > 0\}. \tag{33}$$

The induced map $\mathcal{M}_{\mathcal{N}} \to \mathcal{M}_{\mathcal{E}}$ is surjective and has no critical points. The composition $\mathcal{P}(V) \times \mathcal{P}(U) \to \mathcal{M}_{\mathcal{N}} \to \mathcal{M}_{\mathcal{E}}$ admits a section, thus, we may treat $\mathcal{M}_{\mathcal{E}}$ as

$$\mathcal{M}_{\mathcal{E}} \cong \{(\mathbf{P}, \mathbf{Q}) \mid \|\mathbf{P}\|_2 = \|\mathbf{Q}\|_2 = 1\} \subset \mathcal{P}(\mathbb{R}^p) \times \mathcal{P}(\mathbb{R}^q), \tag{34}$$

which provides $\mathcal{M}_{\mathcal{E}}$ with a smooth manifold structure. The reason we still use $\mathcal{M}_{\mathcal{N}}$ is the metric it possesses, whereas we cannot provide $\mathcal{M}_{\mathcal{E}}$ with a similar metric. Below we demonstrate that the same changes of parametrization and

metric as we utilized in the Gaussian case, make $\widetilde{f}_{\mathcal{E}}$ convex and significantly simplify the treatment. On the other hand, we need $\mathcal{M}_{\mathcal{E}}$ when we talk about the uniqueness of the extremum, since there is no uniqueness of the minimum of $\widetilde{f}_{\mathcal{E}}$ over $\mathcal{M}_{\mathcal{N}}$ due to the scaling ambiguity.

Minimization of $\widetilde{f}_{\mathcal{E}}$ w.r.t. $\mathbf{P}$ and $\mathbf{Q}$ yields a critical point defined by the following system

$$\begin{cases} \mathbf{P} = \dfrac{1}{qn}\displaystyle\sum_{i=1}^{n} \dfrac{\mathbf{X}_i\mathbf{Q}^{-1}\mathbf{X}_i^\top}{\operatorname{Tr}\left(\mathbf{P}^{-1}\mathbf{X}_i\mathbf{Q}^{-1}\mathbf{X}_i^\top\right)}, \\ \mathbf{Q} = \dfrac{1}{pn}\displaystyle\sum_{i=1}^{n} \dfrac{\mathbf{X}_i^\top\mathbf{P}^{-1}\mathbf{X}_i}{\operatorname{Tr}\left(\mathbf{P}^{-1}\mathbf{X}_i\mathbf{Q}^{-1}\mathbf{X}_i^\top\right)}. \end{cases} \tag{35}$$

Similarly to the Gaussian case, there does not exist a closed form solution to this system, and an iterative solution is required which we call the Robust Flip-Flop (RFF). It is also a descent algorithm and converges starting from any initial point due to a similar reasoning. If one wants to remain inside the set $\mathcal{M}_{\mathcal{E}}$ on each iteration, he has to normalize the iterates on each step

$$\begin{cases} \widetilde{\mathbf{P}}_{j+1} = \dfrac{1}{qn}\displaystyle\sum_{i=1}^{n} \dfrac{\mathbf{X}_i\mathbf{Q}_j^{-1}\mathbf{X}_i^\top}{\operatorname{Tr}\left(\mathbf{P}_j^{-1}\mathbf{X}_i\mathbf{Q}_j^{-1}\mathbf{X}_i^\top\right)}, \\ \widetilde{\mathbf{Q}}_{j+1} = \dfrac{1}{pn}\displaystyle\sum_{i=1}^{n} \dfrac{\mathbf{X}_i^\top\mathbf{P}_j^{-1}\mathbf{X}_i}{\operatorname{Tr}\left(\mathbf{P}_j^{-1}\mathbf{X}_i\mathbf{Q}_j^{-1}\mathbf{X}_i^\top\right)}, \\ \mathbf{P}_{j+1} = \dfrac{\widetilde{\mathbf{P}}_{j+1}}{\|\widetilde{\mathbf{P}}_{j+1}\|_2}, \\ \mathbf{Q}_{j+1} = \dfrac{\widetilde{\mathbf{Q}}_{j+1}}{\|\widetilde{\mathbf{Q}}_{j+1}\|_2}. \end{cases} \tag{36}$$

The above reasoning regarding the scaling invariance of the solution explains that when the solution exists and is unique, such normalization does not affect the convergence.

5.3. The Main Statement

In the robust setting described above, a more intuitive result concerning the R-CARMEL estimate and the RFF can be obtained.

Theorem 2. *Assume $X = \{\mathbf{X}_1, \dots \mathbf{X}_n\}$ are independently sampled from a continuous distribution over $\mathcal{M}_{p\times q}$ and consider the problem of minimizing $\widetilde{f}_{\mathcal{E}}(\mathbf{P}, \mathbf{Q}; X)$ over $\mathcal{M}_{\mathcal{E}}$, then*

1. *if* $n < \max[p/q, q/p]$, *there is no unique minimum,*

2. *if* $n > \max[p/q, q/p]$, *there is a unique minimum a.s.,*

3. *if* $n > \max[p/q, q/p]$, *the normalized RFF scheme (36) converges starting from any point of* $\mathcal{M}_{\mathcal{E}}$ *to this unique minimum a.s.*

Proof. The proof can be found in Section 7. □

A few points are in place here.

Remark 4. *Note that unlike the Gaussian case treated before, here it is natural to assume the mean to be known to get a tractable MLE program. This explains the reduction of the necessary number of samples by one in item 1) of Theorem 2 compared to Theorem 1.*

Remark 5. *It is also worth noting that in the robust case the gap between items 1) and 2) consists of one point only, which distinguishes this case from the Gaussian scenario. The robust result clearly provides a sharper threshold between the mode of existence and uniqueness of the MLE and the mode where it does not at all exist.*

6. Proof of Theorem 1

This section treats the Gaussian setting and utilizes a few useful concepts and techniques from commutative algebra. Therefore, for the reader's convenience we transition to a more general treatment of linear spaces, their tensor products, operators over them and a few more related notions. For this purpose, the next section introduces some additional notations.

6.1. Additional Notations

Abstract vector spaces are denoted by capitals V and are assumed to be real Euclidean spaces, their dual spaces are denoted by V^*. Scalar products are denoted by $(\cdot, \cdot)$ and the corresponding norms by $\|\cdot\|$. Spectral norms of operator are denoted by $\|\cdot\|_2$. For an operator $\mathbf{A}$, its adjoint is denoted by $\mathbf{A}^*$. Note that the scalar product induces a canonical isomorphism $V \cong V^*$. If we identify V with the space of columns $\mathbb{R}^p$ in some orthonormal basis, then V^* may be identified with the space of rows $\mathbb{R}^p$. Then the dual basis in V^* is also orthonormal and the isomorphism between V and V^* is given by

the transpose map. For a self-adjoint operator $\mathbf{A}$, we naturally define its real powers and continuous functions of it via continuous functional calculus.

We may naturally identify $V \otimes V^*$ with $\mathrm{End}_{\mathbb{R}}(V)$ via $(\mathbf{v} \otimes \xi)(\mathbf{u}) = \xi(\mathbf{u})\mathbf{v}$, then the scalar product on $V \otimes V^*$ induces a scalar product on $\mathrm{End}_{\mathbb{R}}(V)$ such that for any operators $\mathbf{A}$ and $\mathbf{B}$, $(\mathbf{A}, \mathbf{B}) = \mathrm{Tr}\,(\mathbf{A}\mathbf{B}^*)$. If we identify V with $\mathbb{R}^p$, then $\mathrm{End}_{\mathbb{R}}(V)$ is identified with $\mathcal{M}_{p\times p}(\mathbb{R})$, $\mathbf{A}^*$ becomes $\mathbf{A}^\top$, and $(\mathbf{A}, \mathbf{B}) = \mathrm{Tr}\left(\mathbf{A}\mathbf{B}^\top\right)$. Given two Euclidean spaces V and U, the scalar products on V and U induce one on their tensor product $V \otimes U$.

For any topological space X, we denote its one-point compactification by $\dot{X}$, i.e. $\dot{X} = X \sqcup \{\infty\}$ with the base of neighborhoods of ∞ consisting of the sets $X \setminus K \sqcup \{\infty\}$ for all compact $K \subseteq X$. For a non-compact topological space X, we denote $\{\infty\}$ by ∂X. As an exception, for the real line $\mathbb{R}$, $\dot{\mathbb{R}}$ will be a two point compactification, i.e. $\dot{\mathbb{R}} = \{-\infty\} \sqcup \mathbb{R} \sqcup \{+\infty\}$ endowed with the usual topology making it homeomorphic to the closed unit interval. Given two sequences ω_n and τ_n, we write $\omega_n \asymp \tau_n$ if $\omega_n/\tau_n \to 1$ as $n \to \infty$, while we will usually suppress the n index.

In this section we shall treat the negative log-likelihoods as functions of the inverse matrices $\mathbf{P}^{-1}$ and $\mathbf{Q}^{-1}$ (as we already did while discussing the example in Section 4.3). We do so to simplify calculations and note that this change does affect existence and uniqueness of the extrema in the problem at hand. Therefore, we denote

$$\begin{aligned} f_{\mathcal{N}}(\mathbf{M}, \mathbf{P} \otimes \mathbf{Q}; X) &= \widetilde{f}_{\mathcal{N}}(\mathbf{M}, \mathbf{P}^{-1} \otimes \mathbf{Q}^{-1}; X) \\ &= \frac{1}{n} \sum_{i=1}^{n} \mathrm{Tr}\left(\mathbf{P}(\mathbf{X}_i - \mathbf{M})\mathbf{Q}(\mathbf{X}_i - \mathbf{M})^\top\right) - \ln|\mathbf{P} \otimes \mathbf{Q}|, \quad (37) \end{aligned}$$

$$g_{\mathcal{N}}(\mathbf{P} \otimes \mathbf{Q}; X) = f_{\mathcal{N}}(0, \mathbf{P} \otimes \mathbf{Q}; X) = \frac{1}{n} \sum_{i=1}^{n} \mathrm{Tr}\left(\mathbf{P}\mathbf{X}_i\mathbf{Q}\mathbf{X}_i^\top\right) - \ln|\mathbf{P} \otimes \mathbf{Q}|. \quad (38)$$

6.2. Metric over $\mathcal{P}(V)$

We endow the cone of the positive definite operators over V, $\mathcal{P}(V)$, with a Riemannian metric, whose geodesic connecting any two points $\mathbf{P}, \mathbf{R} \in \mathcal{P}(V)$ is given by

$$\gamma_t(\mathbf{P}, \mathbf{R}) = \mathbf{P}^{\frac{1}{2}} \left(\mathbf{P}^{-\frac{1}{2}}\mathbf{R}\mathbf{P}^{-\frac{1}{2}}\right)^t \mathbf{P}^{\frac{1}{2}}, \quad 0 \leqslant t \leqslant 1. \quad (39)$$

Due to the limited space, we omit a discussion of this metric and its properties. For more details on the relation of this metric to the MLE problems considered here, consult Rapcsák (1991); Wiesel (2011, 2012a,b) and references therein. If we allow t to run over $\mathbb{R}$, then we call the obtained curve an extended geodesic curve.

Fact 1. *A direct computation shows that the Riemannian metric we have just introduced is invariant under inversion. In addition, we note that the log-determinant function is a linear function of t on the geodesic curves $\gamma_t(\mathbf{P}, \mathbf{R})$.*

Lemma 1. *Let V be a vector space, $\mathbf{x} \in V$, $\mathbf{S} \in \mathcal{P}(V)$ and*

$$\varphi_{\mathbf{x}}(t) = (\mathbf{S}^t\mathbf{x}, \mathbf{x}), \quad t \in \mathbb{R}, \tag{40}$$

then its second derivative reads as

$$\varphi''_{\mathbf{x}}(t) = (\ln^2(\mathbf{S})\,\mathbf{S}^t\mathbf{x}, \mathbf{x}) = \left\|\ln(\mathbf{S})\,\mathbf{S}^{\frac{t}{2}}\mathbf{x}\right\|^2 \geqslant 0, \tag{41}$$

in particular, $\varphi_{\mathbf{x}}(t)$ is convex. In addition, the following are equivalent:

1. $\varphi''_{\mathbf{x}}(t) = 0$ *for some* $t \in \mathbb{R}$,
2. $\varphi''_{\mathbf{x}}(t) \equiv 0$ *for all* $t \in \mathbb{R}$,
3. $\mathbf{S}\mathbf{x} = \mathbf{x}$.

Therefore, if $\varphi_{\mathbf{x}}$ is linear in an open neighborhood of some t_0, then $\varphi_{\mathbf{x}}$ is constant on the whole $\mathbb{R}$.

Proof. Since $(\mathbf{S}^t)' = \ln(\mathbf{S})\,\mathbf{S}^t$ and $\ln(\mathbf{S})$ commutes with any power of $\mathbf{S}$, we get (41) and the convexity follows.

Note that 2) implies 1) and, therefore, it is enough to show 1) $\to$ 3) $\to$ 2). If $\mathbf{S}\mathbf{x} = \mathbf{x}$, then for any real t, $\mathbf{S}^t\mathbf{x} = \mathbf{x}$. Hence, $\varphi_{\mathbf{x}}$ is constant on $\mathbb{R}$ and we get 3) $\to$ 2). Now let $\varphi''_x(t) = 0$, then $\ln(\mathbf{S})\,\mathbf{S}^{\frac{t}{2}}\mathbf{x} = 0$. Since $\mathbf{S}$ is invertible, so is $\mathbf{S}^{\frac{t}{2}}$, and $\ln(\mathbf{S})\mathbf{x} = 0$ or, equivalently, $\mathbf{x}$ is an eigenvector of $\mathbf{S}$ with the eigenvalue 1. Finally, the last claim follows from the fact that $\varphi_{\mathbf{x}}(t) = \left\|\mathbf{S}^{\frac{t}{2}}\mathbf{x}\right\|^2 \geqslant 0$ and the only linear nonnegative function is a constant function. □

Lemma 2. *Let V be a vector space, $\mathbf{v} \in V$, and $\mathbf{P}, \mathbf{R} \in \mathcal{P}(V)$, then*

$$\omega_{\mathbf{v}}(t) = (\mathbf{v}, \gamma_t(\mathbf{P}, \mathbf{R})\mathbf{v}) \tag{42}$$

is convex and the following are equivalent:

1. $\omega_{\mathbf{v}}$ *is linear on some open subset of* $\mathbb{R}$,

2. $\omega_{\mathbf{v}}$ *is constant on the whole* $\mathbb{R}$,

3. $\mathbf{Pv} = \mathbf{Rv}$.

Proof. It is an immediate corollary of Lemma 1 if we set $\mathbf{S} = \mathbf{P}^{-\frac{1}{2}}\mathbf{R}\mathbf{P}^{-\frac{1}{2}}$ and $\mathbf{x} = \mathbf{P}^{\frac{1}{2}}\mathbf{v}$. □

6.3. Convexity of $\mathcal{M}_{\mathcal{N}}$ *and* $g_{\mathcal{N}}$

Let V and U be vector spaces, then their tensor product naturally induces a map

$$\otimes\colon \mathcal{P}(V) \times \mathcal{P}(U) \to \mathcal{P}(V \otimes U), \tag{43}$$

sending a pair $(\mathbf{P}, \mathbf{Q})$ to the product $\mathbf{P} \otimes \mathbf{Q}$. We denote the image of this map by $\mathcal{M}_{\mathcal{N}}$. The identification

$$\mathcal{M}_{\mathcal{N}} \cong \{(\mathbf{P}, \mathbf{Q}) \mid \|\mathbf{P}\|_2 = 1\} \subset \mathcal{P}(V) \times \mathcal{P}(U) \tag{44}$$

provides $\mathcal{M}_{\mathcal{N}}$ with a structure of a smooth manifold. Intuitively, this amounts to saying that fixing the norm of the first component of a KP resolves the scaling ambiguity and provides a bijective correspondence between the factors and their products. Note that the normalization in (44) is chosen arbitrarily, and the specific choice does not affect the existence and uniqueness results. In addition, we have

Lemma 3. *The manifold* $\mathcal{M}_{\mathcal{N}} \subset \mathcal{P}(V \otimes U)$ *is convex w.r.t. the geodesic metric defined in Section 6.2.*

Proof. Since $\mathbf{P} \otimes \mathbf{Q} = \mathbf{P} \otimes \mathbf{I} \cdot \mathbf{I} \otimes \mathbf{Q}$ and $\mathbf{P} \otimes \mathbf{I}$ commutes with $\mathbf{I} \otimes \mathbf{Q}$, we have

$$\gamma_t(\mathbf{P} \otimes \mathbf{Q}, \mathbf{R} \otimes \mathbf{T}) = \gamma_t(\mathbf{P}, \mathbf{R}) \otimes \gamma_t(\mathbf{Q}, \mathbf{T}), \tag{45}$$

where the right-hand side is in $\mathcal{M}_{\mathcal{N}}$, thus we are done. □

Below we also make use of the following simple

Fact 2. *For any distinct* $\mathbf{P}, \mathbf{R} \in \mathcal{P}(V)$, $\gamma_t(\mathbf{P}, \mathbf{Q}) \to \infty$ *as* $t \to \pm\infty$.

Lemma 4. *Let* V *be a vector space,* $X \subset V$ *- a fixed finite subset,* $\mathbf{P} \in \mathcal{P}(V)$ *and*

$$g(\mathbf{P}; X) = \frac{1}{|X|} \sum_{\mathbf{x} \in X} (\mathbf{P}\mathbf{x}, \mathbf{x}) - \ln |\mathbf{P}|, \tag{46}$$

then

1. *g is convex w.r.t. the Riemannian metric (39),*

2. *g is linear on $\gamma_t(\mathbf{P},\mathbf{R})$ for some $\mathbf{P},\mathbf{R}\in\mathcal{P}(V)$ if and only if $\mathbf{P}\mathbf{x}=\mathbf{R}\mathbf{x}$ for all $\mathbf{x}\in X$,*

3. *if g has two minima $\mathbf{P}\neq\mathbf{R}$ in $\mathcal{P}(V)$, then the whole extended geodesic $\gamma_t(\mathbf{P},\mathbf{R})$, $t\in\mathbb{R}$ consists of different minima,*

4. *let U be another vector space, $\mathcal{M}_{\mathcal{N}}\subset\mathcal{P}(V\otimes U)$ as before, and $\dot{g}_{\mathcal{N}}\colon\dot{\mathcal{M}}_{\mathcal{N}}\to\dot{\mathbb{R}}$ extends $g_{\mathcal{N}}$ such that $\dot{g}_{\mathcal{N}}(\infty;X)=+\infty$, then $\dot{g}_{\mathcal{N}}$ is continuous if and only if $g_{\mathcal{N}}$ has a unique minimum.*

Proof. 1) $\ln|\mathbf{P}|$ is linear on geodesics and, thus, convex by Fact 1. Lemma 2 implies that each $(\mathbf{P}\mathbf{x},\mathbf{x})$ is convex, therefore, so is $g(\mathbf{P};X)$.

2) This follows from the 3) $\to$ 2) implication of Lemma 2.

3) The convexity implies that the restriction of g onto $\gamma_t(\mathbf{P},\mathbf{R})$, $0\leqslant t\leqslant 1$, is constant and, therefore, linear. Now the 1) $\to$ 2) implication of Lemma 2 finishes the proof.

4) Let us show the sufficiency of the condition. Indeed, if $\dot{g}_{\mathcal{N}}$ is continuous then it achieves a minimum at some interior point (that is the existence). If such minimum is not unique, then by 3), $\dot{g}_{\mathcal{N}}$ must be constant on the whole extended geodesic and cannot be continuous when approaching the boundary, since $\dot{g}_{\mathcal{N}}(\infty;X)=+\infty$.

We proceed to the necessity. Let $\mathbf{S}_0\in\mathcal{M}_{\mathcal{N}}$ be the unique minimum and $\nu=g_{\mathcal{N}}(\mathbf{S}_0;X)$. We denote by $\mathcal{T}_{\mathbf{S}_0}\mathcal{M}_{\mathcal{N}}$ the tangent space to our manifold at the point $\mathbf{S}_0$, and choose a tangent vector $\mathbf{R}\in\mathcal{T}_{\mathbf{S}_0}\mathcal{M}_{\mathcal{N}}$. Let $\gamma_t(\mathbf{R})$ be the geodesic starting at $\mathbf{S}_0$ in direction $\mathbf{R}$,

$$\gamma_0(\mathbf{R})=\mathbf{S}_0\quad\text{and}\quad\gamma_0'(\mathbf{R})=\mathbf{R}. \tag{47}$$

The explicit formula for γ_t reads as

$$\gamma_t(\mathbf{R})=\mathbf{C}_0e^{t\mathbf{C}_0^{-1}\mathbf{R}\mathbf{C}_0^{-1}}\mathbf{C}_0, \tag{48}$$

where $\mathbf{C}_0=\mathbf{S}_0^{\frac{1}{2}}$. In particular, $\gamma_{\lambda t}(\mathbf{R})=\gamma_t(\lambda\mathbf{R})$ for any $\lambda>0$. Set $\delta_t(\mathbf{R})=g_{\mathcal{N}}(\gamma_t(\mathbf{R});X)$. We claim that

$$\sigma=\min_{|\mathbf{R}|=1}\delta_1'(\mathbf{R})>0. \tag{49}$$

Indeed, suppose $\delta_1'(\mathbf{R}_0) = 0$ for some $\mathbf{R}_0$, then since $\delta_t(\mathbf{R}_0)$ is convex and $t = 0$ is the minimum, $\delta_t(\mathbf{R}_0)$ is constant for $0 \leqslant t \leqslant 1$. Thus, the minimum is not unique, which is a contradiction.

Denote by B_t^0 the centered open ball of radius t in $\mathcal{T}_{\mathbf{S}_0}\mathcal{M}_{\mathcal{N}}$ and B_t - its closure, then $\gamma_1(B_t^0)$ is a family of open neighborhoods of $\mathbf{S}_0$ with compact closure $K_t = \gamma_1(B_t)$. Thus, we need to show that

$$\inf_{\mathbf{P}\otimes\mathbf{Q}\notin K_t} g_{\mathcal{N}}(\mathbf{P}\otimes\mathbf{Q}; X) \to +\infty, \text{ as } t \to +\infty. \tag{50}$$

Indeed,

$$\begin{aligned}\inf_{\mathbf{P}\otimes\mathbf{Q}\notin K_t} g_{\mathcal{N}}(\mathbf{P}\otimes\mathbf{Q}; X) &= \inf_{\substack{\mathbf{R}\in\mathcal{T}_{\mathbf{S}_0}\mathcal{M}_{\mathcal{N}},\\ \|\mathbf{R}\|\geqslant t}} g_{\mathcal{N}}(\gamma_1(\mathbf{R}); X)\\ &\geqslant g_{\mathcal{N}}\left(\gamma_{\|\mathbf{R}\|}\left(\frac{\mathbf{R}}{\|\mathbf{R}\|}\right); X\right) \geqslant \sigma(\|\mathbf{R}\| - 1) + \nu \geqslant \sigma(t-1) + \nu, \end{aligned} \tag{51}$$

where in the last line $\mathbf{R} \in \mathcal{T}_{\mathbf{S}_0}\mathcal{M}_{\mathcal{N}}$ is any matrix of norm at least t. □

6.4. The Set of "Bad" Samples

Depending on the set X, $g_{\mathcal{N}}$ may happen to be not strictly convex on $\mathcal{M}_{\mathcal{N}}$, or equivalently, $\dot{g}_{\mathcal{N}}$ is not necessarily continuous. In this section we discuss when this situation occurs and compute the measure of the set of samples making $\dot{g}_{\mathcal{N}}$ discontinuous.

For a vector space V of dimension p and natural numbers $d, s \in \mathbb{N}$, define

$$\mathrm{G}_{d\,s}(V) = \{(\mathbf{v}_1, \ldots, \mathbf{v}_s) \in V^s \mid \dim\langle \mathbf{v}_1, \ldots, \mathbf{v}_s\rangle = d\} \subseteq V^s \tag{52}$$

to be the set of all s-tuples of vectors in V, spanning subspaces of dimension d. $\mathrm{G}_{d\,s}(V)$ is a smooth manifold of dimension $(p + s - d)d$. Note also that $\mathrm{G}_{d\,d}(V) \subseteq V^d$ is an open subset, moreover, if we represent $V^d = V \otimes \mathbb{R}^d$, we get an action of $\mathrm{GL}_d(\mathbb{R})$ on V^d, which restricts correctly onto $\mathrm{G}_{d\,d}(V)$ and is free.

Before giving a precise statement about what Diagram 1 displays, let us provide an intuitive explanation of it. The operator analog of a $p \times q$ matrix is an element of $\mathrm{Hom}(U, V)$, thus our n matrix measurements in X together represent an element of $\mathrm{Hom}(U, V)^n$. We next take d linearly independent vectors in a q dimensional U and apply all the elements of our n-tuple to them - this provides us with dn vectors in a p dimensional V. If we now

$$\begin{array}{ccccc} & & \mathcal{A}_{n\,d} = \mathrm{Hom}(U,V)^n \times \mathrm{G}_{d\,d}(U) & \xrightarrow{\Psi} & V^{dn} \\ & & \cup & & \cup \\ \mathrm{Hom}(U,V)^n & \xleftarrow{\pi} & \mathcal{Z}_{d\,n\,r} & \longrightarrow & \mathrm{G}_{r\,dn}(V) \end{array}$$

Diag. 1: Diagram from Definition 2.

take all the pairs of such n-tuples of operators and d-tuples of vectors in U, and consider the described action of the former on the latter, which we call Ψ, we get the first line of the diagram at hand. Let us now consider all the sets of dn vectors inside V spanning subspaces of dimension r and take their preimage under Ψ. We get a subset $\mathcal{Z}_{d\,n\,r}$ of $\mathrm{Hom}(U,V)^n \times \mathrm{G}_{d\,d}(U)$ depicted in the diagram. Finally, the leftmost arrow π denotes the projection of this set onto $\mathrm{Hom}(U,V)^n$. This informal description is made precise by the following

Definition 2. *Let V and U be vector spaces of dimensions p and q, respectively, and $n, d, r \in \mathbb{N}$ be such that $d \leqslant q$. Consider Diagram 1, where Ψ is defined as*

$$\Psi : \mathrm{Hom}(U,V)^n \times \mathrm{G}_{d\,d}(U) \to V^{dn}, \tag{53}$$

$$((\varphi_1, \ldots, \varphi_n), (\mathbf{u}_1, \ldots, \mathbf{u}_d)) \mapsto (\varphi_i \mathbf{u}_j). \tag{54}$$

Identify $V^{dn} = V^n \otimes \mathbb{R}^d$, then Ψ reads as

$$\Psi((\varphi_1, \ldots, \varphi_n), (\mathbf{u}_1, \ldots, \mathbf{u}_d)) = \varphi_i \mathbf{u}_j \otimes \mathbf{e}_j, \tag{55}$$

and $\mathbf{C} \in \mathrm{GL}_d(\mathbb{R})$ acts by $\mathbf{I} \otimes \mathbf{C}$. Define $\mathcal{Z}_{d\,n\,r} = \Psi^{-1}\left(\mathrm{G}_{r\,dn}(V)\right)$ and π to be the restriction of the projection along $\mathrm{G}_{d\,d}(U)$.

Lemma 5. *With the notations of Definition 2, for each $(\mathbf{u}_1, \ldots, \mathbf{u}_d) \in \mathrm{G}_{d\,d}(U)$ the map*

$$\Omega : \mathrm{Hom}(U,V)^n \to V^{dn}, \tag{56}$$

$$(\varphi_1, \ldots, \varphi_n) \mapsto (\varphi_i \mathbf{u}_j) \tag{57}$$

is surjective.

Proof. Ω is a direct sum of n maps

$$\omega\colon\ \mathrm{Hom}(U,V) \to V^d, \tag{58}$$

$$\varphi \mapsto (\varphi \mathbf{u}_1, \ldots, \varphi \mathbf{u}_d). \tag{59}$$

Now we choose bases in V and U such that $V = \mathbb{R}^p$ and $U = \mathbb{R}^q$. Then $\operatorname{Hom}(U, V) = \mathcal{M}_{p\times q}$ and $V^d = \mathcal{M}_{p\times d}$. Let $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_d]$, then ω reads as

$$\omega\colon \mathcal{M}_{p\times q}(\mathbb{R}) \to \mathcal{M}_{p\times d}(\mathbb{R}), \tag{60}$$

$$\mathbf{X} \mapsto \mathbf{X}\mathbf{U}. \tag{61}$$

Since $d \leqslant q$ and the columns of $\mathbf{U}$ are linearly independent, $\mathbf{U}$ is of full rank. Hence, the map ω is surjective.

□

Lemma 6. *With the notations of Definition 2,*

1. *Ψ is surjective and $d\Psi$ is surjective at each point,*
2. *$\mathcal{Z}_{dnr}$ is a smooth manifold with*

$$\dim \mathcal{Z}_{dnr} = pqn + qd - (p - r)(dn - r), \tag{62}$$

3. *a non empty fiber of π has dimension at least d^2.*

Proof. 1) To show Ψ is surjective, let $(\mathbf{u}_1, \dots, \mathbf{u}_d) \in \mathrm{G}_{d\,d}(U)$. Now it is enough to show that

$$\Psi(\cdot, (\mathbf{u}_1, \dots, \mathbf{u}_d))\colon \operatorname{Hom}(U, V)^n \to V^{dn} \tag{63}$$

is surjective, which follows from Lemma 5. We proceed to the surjectivity of $d\Psi$ at any point $(\{\varphi_i\}, \{\mathbf{u}_j\})$. Since $\mathrm{G}_{d\,d}(U) \subseteq U^d$ is open, identify the tangent space of $\mathrm{G}_{d\,d}(U)$ at $\{\mathbf{u}_j\}$ with U^d, then

$$d_{(\{\varphi_i\},\{\mathbf{u}_j\})}\Psi\colon \operatorname{Hom}(U, V)^n \times U^d \to V^{dn}, \tag{64}$$

$$(\{\varphi_i^*\}, \{\mathbf{u}_j^*\}) \mapsto (\varphi_i^*\mathbf{u}_j + \varphi_i\mathbf{u}_j^*). \tag{65}$$

Taking $\mathbf{u}_j^* = 0$, $j = 1, \dots, d$, we get $(\{\varphi_i^*\}, \{\mathbf{0}\}) \mapsto (\varphi_i^*\mathbf{u}_j)$ and the result follows from Lemma 5.

2) The fact that $\mathcal{Z}_{dnr}$ is a smooth manifold follows from 1) and the Implicit Function Theorem. A direct computation yields

$$\begin{aligned}\dim \mathcal{Z}_{dnr} &= \dim \mathcal{A}_{n\,d} + \dim \mathrm{G}_{r\,dn}(V) - \dim V^{dn} \\ &= pqn + qd + (p + dn - r)r - pdn = pqn + qd - (p - r)(dn - r).\end{aligned} \tag{66}$$

3) $\mathrm{GL}_d(\mathbb{R})$ acts freely on the fibers of π, so the dimension of a fiber is at least $\dim \mathrm{GL}_d(\mathbb{R}) = d^2$. □

Corollary 1. *With the notations of Lemma 6, if $n > p/q+q/p$ and $r \leqslant dp/q$, then $\pi(\mathcal{Z}_{dnr})$ is a set of Lebesgue measure zero.*

Proof. By Sard's theorem, it is enough to show that the image of π consists of critical values only. So we need to show that

$$\operatorname{rk}(\pi) < \dim \operatorname{Hom}(U,V)^n. \tag{67}$$

Since

$$\operatorname{rk}(\pi) \leqslant \dim \pi(\mathcal{Z}_{dnr}) - d^2, \tag{68}$$

by Lemma 6 item 3), it is enough to prove that

$$\dim \pi(\mathcal{Z}_{dnr}) - d^2 < \dim \operatorname{Hom}(U,V)^n. \tag{69}$$

By Lemma 6 item 2), the latter is equivalent to

$$pqn + qd - (p-r)(dn-r) < pqn, \tag{70}$$

therefore, we need

$$\frac{q-d}{p-r} + \frac{r}{d} < n, \quad \forall\, 1 \leqslant d \leqslant q \text{ and } 0 \leqslant r \leqslant d\frac{p}{q}. \tag{71}$$

Differentiate the left hand-side w.r.t. r to see that it is a strictly increasing function of r, thus, it is enough to demonstrate the inequality for $r = dp/q$, which is

$$\frac{q-d}{p-d\frac{p}{q}} + \frac{p}{q} = \frac{q}{p} + \frac{p}{q} < n, \tag{72}$$

and holds by the assumption. □

Corollary 2. *Let V and U be vector spaces of dimensions p and q, respectively, and X - a finite mutually continuous family of random operators from U to V such that*

$$|X| > \frac{p}{q} + \frac{q}{p}, \tag{73}$$

then for each random subspace $E \subseteq V$ we have

$$\frac{\dim \sum_{\mathbf{X}\in X} \mathbf{X}E}{\dim E} > \frac{p}{q} \quad a.s. \tag{74}$$

Proof. Note that X is distributed over $\operatorname{Hom}(U,V)^n$, then

$$\left\{X \;\middle|\; \exists E \subseteq U:\ \frac{\dim \sum_{\mathbf{X}\in X} \mathbf{X}E}{\dim E} \leqslant \frac{p}{q}\right\} = \bigcup_{\substack{1\leqslant d\leqslant q\\ 0\leqslant r\leqslant d\frac{p}{q}}} \pi(\mathcal{Z}_{d\,n\,r}), \tag{75}$$

and the result follows from Corollary 1. □

Note that in Corollary 2, we do not care whether $p \geqslant q$ or $q \geqslant p$. When $p \geqslant q$, the inequality may fail if $\mathbf{X}E$ is not big enough compared to E, and when $q \geqslant p$, it may fail if E belongs to the kernels of all samples $\mathbf{X}$.

6.5. Flags

In the proof of the main theorem (see Theorem 3 below) we will analyze the behavior of $g_{\mathcal{N}}(\mathbf{P}\otimes\mathbf{Q}; X)$ when $\mathbf{P}\otimes\mathbf{Q}$ tends to ∞. There are many ways $\mathbf{P}\otimes\mathbf{Q}$ may tend to the boundary and in order to classify all possibilities we utilize the flag machinery introduced next.

Definition 3. *Let U be a vector space, then a flag $\mathcal{F}$ of length s in U is an ascending sequence of proper subspaces of U*

$$\mathcal{F} = \{0 = U_0 \subsetneq U_1 \subsetneq \ldots \subsetneq U_s \mid U_s \subseteq U\}. \tag{76}$$

The flag is called non-trivial if $0 \subsetneq U_1 \subsetneq U$. A subsequence of $\mathcal{F}$ is called a subflag. Let V be another vector space and $\mathcal{G} = \{V_i\}$ be a flag of length s in V. Let $\zeta\colon U \to V$ be a linear map with $\zeta U_i \subseteq V_i$ for each $i \leqslant s$, then we write $\zeta\mathcal{F} \subseteq \mathcal{G}$. In addition, for all $0 \leqslant i, j \leqslant s$ define

$$\Pi(\mathcal{F},\mathcal{G})_{i\,j} = q(\dim V_j - \dim V_i) - p(\dim U_j - \dim U_i). \tag{77}$$

If $\mathbf{r} = \{r_1 > \ldots > r_s > 0\}$ is a vector of strictly decreasing real numbers, we define

$$C(\mathcal{F},\mathcal{G},\mathbf{r}) = \sum_{i=1}^{s} r_i \Pi(\mathcal{F},\mathcal{G})_{i-1\,i}, \tag{78}$$

and note that

$$\Pi(\mathcal{F},\mathcal{G})_{i\,j} + \Pi(\mathcal{F},\mathcal{G})_{j\,k} = \Pi(\mathcal{F},\mathcal{G})_{i\,k}. \tag{79}$$

Lemma 7. *Let V, U and $\mathcal{G}, \mathcal{F}$ be as above. If $\Pi(\mathcal{F}, \mathcal{G})_{0\,i} > 0$ for $i = 1, \ldots, s$, then there exist subflags $\mathcal{F}' \subseteq \mathcal{F}$, $\mathcal{G}' \subseteq \mathcal{G}$, and a subsequence $\mathbf{r}' \subseteq \mathbf{r}$ - all of length s', such that*

$$\Pi(\mathcal{F}', \mathcal{G}')_{0\,1} > 0, \quad \Pi(\mathcal{F}', \mathcal{G}')_{i-1\,i} \geqslant 0, \;\; i = 1, \ldots, s', \tag{80}$$

and $C(\mathcal{F}, \mathcal{G}, \mathbf{r}) \geqslant C(\mathcal{F}', \mathcal{G}', \mathbf{r}')$. In particular, $C(\mathcal{F}, \mathcal{G}, \mathbf{r}) > 0$.

Proof. The proof is by induction on s and the base $s = 1$ is the hypothesis of the lemma. Now let $s > 1$ and assume the claim fails for $\mathcal{F}$ and $\mathcal{G}$, then for some $t \leqslant s$

$$\Pi(\mathcal{F}, \mathcal{G})_{i-1\,i} \geqslant 0 \quad \forall i < t, \quad \text{and} \quad \Pi(\mathcal{F}, \mathcal{G})_{t-1\,t} < 0. \tag{81}$$

Construct $\mathcal{F}'$ and $\mathcal{G}'$ from $\mathcal{F}$ and $\mathcal{G}$ by excluding U_{t-1} and V_{t-1}, respectively, and $\mathbf{r}'$ from $\mathbf{r}$ by excluding r_{t-1}, then

$$\begin{aligned} C(\mathcal{F}, \mathcal{G}, \mathbf{r}) &= \sum_{i \neq t-1, t} r_i \Pi(\mathcal{F}, \mathcal{G})_{i-1\,i} + r_{t-1} \Pi(\mathcal{F}, \mathcal{G})_{t-2\,t-1} + r_t \Pi(\mathcal{F}, \mathcal{G})_{t-1\,t} \\ &\geqslant \sum_{i \neq t-1, t} r_i \Pi(\mathcal{F}, \mathcal{G})_{i-1\,i} + r_t \left(\Pi(\mathcal{F}, \mathcal{G})_{t-2\,t-1} + \Pi(\mathcal{F}, \mathcal{G})_{t-1\,t} \right) = C(\mathcal{F}', \mathcal{G}', \mathbf{r}'). \end{aligned} \tag{82}$$

□

6.6. The Known Mean Case

In this section, we prove our main result assuming the mean to be known. For this purpose we need a few auxiliary results.

Lemma 8. *Let λ, μ and γ be positive sequences such that $\lambda\mu = O(\ln \gamma)$ and $\gamma \to +\infty$, then*

$$\ln \lambda^{-1} \geqslant \ln \mu + o(\ln \gamma). \tag{83}$$

Proof. $\lambda\mu = \alpha \ln \gamma$, with $\alpha \leqslant \kappa$ for some constant κ. Taking logarithms yields

$$\ln \lambda + \ln \mu = \ln \, \ln \gamma + \ln \alpha, \tag{84}$$

hence,

$$\ln \lambda^{-1} = \ln \mu - \ln \, \ln \gamma - \ln \alpha \geqslant \ln \mu - \ln \, \ln \gamma - \ln \kappa = \ln \mu + o(\ln \gamma). \tag{85}$$

□

Theorem 3. *Let V and U be vector spaces of dimensions p and q, respectively, $\mathcal{M}_{\mathcal{N}} \subset \mathcal{P}(V \otimes U)$ as in (43) and $X \subset V \otimes U$ - a finite mutually continuous family of random vectors such that $|X| > p/q + q/p$, then $g_{\mathcal{N}} \colon \mathcal{M}_{\mathcal{N}} \to \mathbb{R}$ extends to a continuous function $\dot{g}_{\mathcal{N}} \colon \dot{\mathcal{M}}_{\mathcal{N}} \to \dot{\mathbb{R}}$ via $\dot{g}_{\mathcal{N}}(\infty; X) = +\infty$. In particular, there exists a unique minimum of $g_{\mathcal{N}}$ over $\mathcal{M}_{\mathcal{N}}$.*

Remark 6. *Note that the statement allows the members of X to be statistically dependent and does not require identical distribution. This generality is necessary when we treat the case of manually empirically centered samples below and makes application of Theorem 3 possible without additional adjustments.*

Proof. By Lemma 4 item 4) it is enough to show that $g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; X) \to +\infty$ as $(\mathbf{P}, \mathbf{Q}) \to \infty$. Suppose on the contrary, that there exists a sequence $(\mathbf{P}, \mathbf{Q}) \to \infty$ (we omit the j indexing in $\{(\mathbf{P}_j, \mathbf{Q}_j)\}_j$ to simply notations) such that $g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; X) \leqslant \kappa$ for some constant κ. Rewrite $g_{\mathcal{N}}$ as

$$g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; X) = \sum_{\mathbf{X} \in X} \varphi_{\mathbf{X}}(\mathbf{P}, \mathbf{Q}) + \psi(\mathbf{P}, \mathbf{Q}) = \varphi(\mathbf{P}, \mathbf{Q}) + \psi(\mathbf{P}, \mathbf{Q}), \tag{86}$$

where

$$\varphi_{\mathbf{X}}(\mathbf{P}, \mathbf{Q}) = \frac{1}{|X|}(\mathbf{PXQ}, \mathbf{X}), \quad \psi(\mathbf{P}, \mathbf{Q}) = -\ln |\mathbf{P} \otimes \mathbf{Q}|. \tag{87}$$

Recall that $\mathcal{M}_{\mathcal{N}}$ can be identified with

$$\mathcal{M}_{\mathcal{N}} \cong \{(\mathbf{P}, \mathbf{Q}) \mid \|\mathbf{P}\|_2 = 1\} \subset \mathcal{P}(V) \times \mathcal{P}(U). \tag{88}$$

If $\|\mathbf{Q}\|_2$ is bounded, then the sequence $(\mathbf{P}, \mathbf{Q})$ tends to a singular pair (at least one of the matrices tends to a singular limit). In this case, $\varphi(\mathbf{P}, \mathbf{Q}) \asymp O(1)$ and $\psi(\mathbf{P}, \mathbf{Q}) \to +\infty$.

Now assume $\|\mathbf{Q}\|_2 \to +\infty$, the only problem here is that $-p \ln |\mathbf{Q}|$ may tend to $-\infty$. We should show that we can compensate for this with the other summands. Let $\boldsymbol{\sigma}^{\mathbf{Q}}$ be the spectrum of $\mathbf{Q}$, then it can be partitioned as $\boldsymbol{\sigma}^{\mathbf{Q}} = \boldsymbol{\sigma}^{\mathbf{Q}}_{\infty} \sqcup \boldsymbol{\sigma}^{\mathbf{Q}}_{m}$ such that

- for each $\mu \in \boldsymbol{\sigma}^{\mathbf{Q}}_{\infty}$, $\ln \mu \asymp r_\mu \ln \|\mathbf{Q}\|_2$, where $r_\mu > 0$ is constant,
- for each $\mu \in \boldsymbol{\sigma}^{\mathbf{Q}}_{m}$, $\ln \mu = o(\ln \|\mathbf{Q}\|_2)$.

Order the elements of $\boldsymbol{\sigma}_\infty^{\mathbf{Q}}$ by their rate of convergence

$$\boldsymbol{\sigma}_\infty^{\mathbf{Q}} = \boldsymbol{\sigma}_1^{\mathbf{Q}} \sqcup \ldots \sqcup \boldsymbol{\sigma}_s^{\mathbf{Q}}, \tag{89}$$

where

- for each $\mu \in \boldsymbol{\sigma}_1^{\mathbf{Q}},\ \mu \asymp \|\mathbf{Q}\|_2$,
- for each $\mu, \mu' \in \boldsymbol{\sigma}_i^{\mathbf{Q}},\ \lim \mu/\mu'$ is a non-zero constant,
- for any i, if $\mu_i \in \boldsymbol{\sigma}_i^{\mathbf{Q}}$, then $\mu_{i+1} = o(\mu_i)$.

For a fixed i, let $\{\bar{K}_i\}$ be the sequence of random subspaces of U generated by the eigenvectors corresponding to $\boldsymbol{\sigma}_i$, and K_i be the limit of $\{\bar{K}_i\}$ (it exists after passing to an appropriate subsequence, if needed). Now $U_k = \oplus_{j \leqslant k} K_j$ form a non-trivial random flag of length s in U,

$$\mathcal{F} = \{0 = U_0 \subseteq U_1 \subseteq U_2 \subseteq \ldots \subseteq U_s \mid U_s \subseteq U\}. \tag{90}$$

Let $\boldsymbol{\sigma}^{\mathbf{P}}$ be the spectrum of $\mathbf{P}$ and set

$$\boldsymbol{\sigma}_i^{\mathbf{P}} = \{\lambda \in \boldsymbol{\sigma}^{\mathbf{P}} \mid \lambda\mu_i = O(\ln \|\mathbf{Q}\|_2), \text{ for } \mu_i \in \boldsymbol{\sigma}_i^{\mathbf{Q}}\}. \tag{91}$$

By the definition of $\boldsymbol{\sigma}_i^{\mathbf{Q}}$, $\boldsymbol{\sigma}_i^{\mathbf{P}}$ does not depend on the choice of $\mu_i \in \boldsymbol{\sigma}_i^{\mathbf{Q}}$. Let $\{\bar{L}_i\}$ be the sequence of random subspaces of V generated by the eigenvectors corresponding to $\boldsymbol{\sigma}_i^{\mathbf{P}}$ and L_i be the limit of $\{\bar{L}_i\}$ (here again, it exists after passing to an appropriate subsequence). Denote $V_k = \oplus_{j \leqslant k} L_j$ and define

$$\mathcal{G} = \{0 = V_0 \subseteq V_1 \subseteq \ldots \subseteq V_s \mid V_s \subseteq V\}, \tag{92}$$

which is a flag of length s in V. We denote the orthogonal projector in $V \otimes U$ onto $L_j \otimes K_i$ by π_{ji}.

Now there are two possibilities:

- $\exists \mathbf{X} \in X : \mathbf{X}\mathcal{F} \not\subseteq \mathcal{G}$. Let $\mathbf{X}U_i \not\subseteq V_i$, i.e. there is some j such that $\ln \|\mathbf{Q}\|_2 = o(\lambda_i\mu_j)$ and $\pi_{ij}(\mathbf{X}) \neq 0$, then

$$\varphi(\mathbf{P}, \mathbf{Q}) \geqslant \varphi_{\mathbf{X}}(\mathbf{P}, \mathbf{Q}) = \frac{1}{|X|}(\mathbf{PXQ}, \mathbf{X}) \geqslant \frac{1}{|X|}\lambda_i\mu_j|\pi_{ij}(\mathbf{X})|^2, \tag{93}$$

and $\psi(\mathbf{P}, \mathbf{Q}) = O(\ln \|\mathbf{Q}\|_2) = o(\lambda_i\mu_j)$, hence, $g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; X) \to +\infty$.

- $\forall \mathbf{X} \in X : \mathbf{X}\mathcal{F} \subseteq \mathcal{G}$. Since $\varphi(\mathbf{P}, \mathbf{Q}) \not\to -\infty$, we can ignore this summand when considering the asymptotic behavior. Compute $\psi(\mathbf{P}, \mathbf{Q}) = -q \ln |\mathbf{P}| - p \ln |\mathbf{Q}|$ explicitly,

$$-q \ln |\mathbf{P}| = q \sum_{i=1}^{s} (\dim V_i - \dim V_{i-1}) \ln \lambda_i^{-1} - q \ln \det \mathbf{P}|_{V_s^\perp}, \tag{94}$$

$$-p \ln |\mathbf{Q}| = -p \sum_{i=1}^{s} (\dim U_i - \dim U_{i-1}) \ln \mu_i + o(\ln \|\mathbf{Q}\|_2), \tag{95}$$

where $\lambda_i \in \boldsymbol{\sigma}_i^{\mathbf{P}}$ and $\mu_i \in \boldsymbol{\sigma}_i^{\mathbf{Q}}$. Note that $-q \ln \det \mathbf{P}|_{V_s^\perp} \not\to -\infty$, therefore, we may drop this summand. Since $\mu_i \asymp r_i \ln \|\mathbf{Q}\|_2$, by Lemma 8 we obtain

$$-\ln |\mathbf{P} \otimes \mathbf{Q}| \gtrsim C(\mathcal{F}, \mathcal{G}, \mathbf{r}) \ln \|\mathbf{Q}\|_2 + o(\ln \|\mathbf{Q}\|_2), \tag{96}$$

thus, it is enough to prove that the coefficient $C(\mathcal{F}, \mathcal{G}, \mathbf{r})$ is positive. This would follow from Lemma 7 if we prove that $\Pi(\mathcal{F}, \mathcal{G})_{0\,i} > 0$ for $1 \leqslant i \leqslant s$. Indeed,

$$\Pi(\mathcal{F}, \mathcal{G})_{0\,i} = q \dim V_i - p \dim U_i = q \dim U_i \left(\frac{\dim XU_i}{\dim U_i} - \frac{p}{q} \right). \tag{97}$$

Since $\mathcal{F}$ is non-trivial, $\dim U_i \neq 0$ for $i \geqslant 1$. $|X| > p/q + q/p$, thus, due to Corollary 2 the expression in brackets is a.s. positive. This finishes the proof.

□

The proof we have just presented may be complicated to grasp due a large amount of new notations and technical details, therefore, we now explain it in an informal way. Using the same notations, let us describe the main point of using flags. Choose bases in U and V respecting the subspaces K_i and L_j. In these bases all the samples $\mathbf{X} \in X$ has s blocks of rows and s block of columns corresponding to L_i and K_i, respectively. Hence, each sample consists of s^2 blocks. Now we can easily count the contributions of the blocks to the asymptotic of $g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; X)$.

The contributed speed of the (i, j)-th block of any $\mathbf{X}$ is $\lambda_i \mu_j$, up to a scalar depending on $\mathbf{X}$. In order to determine the asymptotic behavior of $g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; X)$ written as in (86), we need to compare the negative impact of

$\psi(\mathbf{P}, \mathbf{Q})$ with the positive one of $\varphi(\mathbf{P}, \mathbf{Q})$. The highest rate negative summand appearing in $\psi(\mathbf{P}, \mathbf{Q})$ decreases with the rate of at most $\ln \|\mathbf{Q}\|_2$ up to a fixed scalar. If $\lambda_i \mu_j$ tends to infinity faster than $\ln \|\mathbf{Q}\|_2$, then $g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; X)$ would tend to $+\infty$.

The problem occurs if all the blocks corresponding to those $\lambda_i \mu_j$ growing faster than $\ln \|\mathbf{Q}\|_2$ are zero for all $\mathbf{X} \in X$. Let us note that if the (i, j)-th block is zero for all $\mathbf{X}$, then all the blocks with smaller j and higher i (to the left and down of our block) have higher speed and, hence, must be zero (otherwise we are in the first situation). This precisely means that all the samples $\mathbf{X} \in X$ are block upper triangular, i.e. they map flag $\mathcal{F}$ into $\mathcal{G}$.

Now we just use these observations together with Lemmas 7 and 8 to explicitly calculate the leading asymptotic term of $\psi(\mathbf{P}, \mathbf{Q})$, which thanks to Corollary 2 grows to $+\infty$ and, therefore, implies the desired.

6.7. The Unknown Mean Case

Let V be a vector space, $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset V$ and let $\{\mathbf{e}_i\} \subset \mathbb{R}^n$ be the standard basis. Define an element $\mathbf{x}^* \in V \otimes \mathbb{R}^n$ as $\mathbf{x}^* = \sum_{i=1}^n \mathbf{x}_i \otimes \mathbf{e}_i$. Then for any $\mathbf{P} \in \mathcal{P}(V)$,

$$\sum_{\mathbf{x} \in X} (\mathbf{P}\mathbf{x}, \mathbf{x}) = (\mathbf{P} \otimes \mathbf{I}\, \mathbf{x}^*, \mathbf{x}^*). \tag{98}$$

Let now introduce the sample mean

$$\widehat{\mathbf{x}} = \frac{1}{|X|} \sum_{\mathbf{x} \in X} \mathbf{x}, \tag{99}$$

and $\mathbf{1} = [1, \dots, 1]^\top \in \mathbb{R}^n$.

Lemma 9.

$$\mathbf{S} = \begin{pmatrix} 1 - \frac{1}{n} & -\frac{1}{n} & \dots & -\frac{1}{n} \\ -\frac{1}{n} & 1 - \frac{1}{n} & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ -\frac{1}{n} & -\frac{1}{n} & \dots & 1 - \frac{1}{n} \end{pmatrix} \in \mathbb{R}^{n \times n} \tag{100}$$

is an orthogonal projector onto a subspace of codimension 1.

Proof. The spectrum of $\mathbf{S}$ contains one 0 eigenvalue and the rest $n-1$ eigenvalues are 1-s. □

Lemma 10. *Let V be a vector space, $X \subset V$ be a finite mutually continuous set of vectors, and $\mathbf{P} \in \mathcal{P}(V)$ - a random operator, then there exists a set $Z \subset V$ of mutually continuous vectors such that $|Z| = |X| - 1$ and*

$$\frac{1}{|X|}\sum_{\mathbf{x}\in X}(\mathbf{P}(\mathbf{x}-\widehat{\mathbf{x}}),\mathbf{x}-\widehat{\mathbf{x}}) = \frac{1}{|Z|}\sum_{\mathbf{z}\in Z}(\mathbf{P}\mathbf{z},\mathbf{z}). \tag{101}$$

Proof. Note that $\mathbf{x}^* - \widehat{\mathbf{x}} \otimes \mathbf{1} = (\mathbf{I} \otimes \mathbf{S})\,\mathbf{x}^*$. Let $\{\mathbf{f}_i\}$ be another orthonormal basis of $\mathbb{R}^n$ such that $\ker \mathbf{S} = \langle \mathbf{f}_n \rangle$. Compute

$$(\mathbf{I} \otimes \mathbf{S})\,\mathbf{x}^* = \sum_{i=1}^{n-1} \mathbf{y}_i \otimes \mathbf{f}_i, \tag{102}$$

and denote $Y = \{\mathbf{y}_1, \ldots, \mathbf{y}_{n-1}\}$, then

$$\sum_{\mathbf{x}\in X}(\mathbf{P}(\mathbf{x}-\widehat{\mathbf{x}}),\mathbf{x}-\widehat{\mathbf{x}}) = ((\mathbf{P}\otimes\mathbf{S})\,\mathbf{x}^*,\mathbf{x}^*) = (\mathbf{P}\otimes\mathbf{I}\,\mathbf{y}^*,\mathbf{y}^*) = \sum_{\mathbf{y}\in Y}(\mathbf{P}\mathbf{y},\mathbf{y}), \tag{103}$$

where $\mathbf{y}_i$ are now centered.

Take

$$\mathbf{z}_i = \sqrt{\frac{|X|-1}{|X|}}\mathbf{y}_i, \quad i = 1, \ldots, n-1. \tag{104}$$

The mutual continuity of $\mathbf{z}_i$ follows from the fact that the function $V^n \to V^{n-1}$ mapping X to Z is linear and surjective. □

Lemma 11. *Let V and U be vector spaces and $X \subset V \otimes U$ - a mutually continuous family of random vectors with $|X| > 1 + p/q + q/p$, then $f_{\mathcal{N}}(\widehat{\mathbf{X}}, \mathbf{P}, \mathbf{Q}; X)$ as a function of $\mathbf{P}$ and $\mathbf{Q}$ extends to a continuous $\dot{f}_{\mathcal{N}} \colon \dot{\mathcal{M}}_{\mathcal{N}} \to \dot{\mathbb{R}}$ via $\dot{f}_{\mathcal{N}}(\infty) = +\infty$, in particular, $f_{\mathcal{N}}$ has a unique minimum over $\mathcal{M}_{\mathcal{N}}$.*

Proof. Applying Lemma 10, we get a mutually continuous set $Z \subset V \otimes U$ such that $|Z| = |X| - 1$ and $f_{\mathcal{N}}(\widehat{\mathbf{X}}, \mathbf{P}, \mathbf{Q}; X) = g_{\mathcal{N}}(\mathbf{P}, \mathbf{Q}; Z)$. Now the desired claim follows from Theorem 3. □

Theorem 4. *Let V and U be vector spaces, $X \subset V \otimes U$ be a mutually continuous family of random vectors with $|X| > 1 + p/q + q/p$, and $S = V \otimes U \times \mathcal{M}_{\mathcal{N}}$, then $f_{\mathcal{N}} \colon S \to \mathbb{R}$ extends to a continuous function $\dot{f}_{\mathcal{N}} \colon \dot{S} \to \dot{\mathbb{R}}$ via $\dot{f}_{\mathcal{N}}(\infty) = +\infty$ and, in particular, there exists a unique minimum of $f_{\mathcal{N}}$ over S.*

Proof. Note that for a fixed pair $(\mathbf{P}, \mathbf{Q})$, the value of $\mathbf{M}$ minimizing $f_{\mathcal{N}}$ is the sample average, which does not depend on the values of $\mathbf{P}$ and $\mathbf{Q}$. Therefore, the result follows from Lemma 11 and Theorem 3. □

7. Proof of Theorem 2

The proof in this section is quite similar to that given in Soloveychik and Wiesel (2016), thus, we made it less verbose than the proof of the previous section. For more details please consult Soloveychik and Wiesel (2016). Analogously to Definition 3, we introduce the notion of a descending flag and note that the usage of flags in this section is different from that of Section 6.

Definition 4. *Let V be a real linear space, $X \subset V$ be a finite subset and $\mathcal{F} = \{V = V_0 \supsetneq V_1 \supsetneq \dots V_s \supsetneq V_s \supseteq 0\}$ be a descending flag of length s in V. Define*

$$\Delta(\mathcal{F}, X)_{i\,j} = \dim V_i - \dim V_j - \frac{\dim V_1}{|X|}\left(|X \cap V_i| - |X \cap V_j|\right), \tag{105}$$

where $0 \leqslant i,\, j \leqslant s$. In addition, given a decreasing sequence

$$\mathbf{r} = \{r_1 > \dots > r_s\} \subset \mathbb{R} \tag{106}$$

of length s, define

$$S(\mathcal{F}, X, \mathbf{r}) = \sum_{i=1}^{s} r_i \Delta(\mathcal{F}, X)_{i-1\,i}. \tag{107}$$

It now follows immediately from the definition that

$$\Delta(\mathcal{F}, X)_{i\,j} + \Delta(\mathcal{F}, X)_{j\,k} = \Delta(\mathcal{F}, X)_{i\,k}, \quad i, j, k = 0, \dots, s. \tag{108}$$

Lemma 12. *Let $X \subseteq V$ be a finite subset, $\mathcal{F}$ be a flag of length s in V, $\mathbf{r}$ be a sequence as in (106), and $\Delta(\mathcal{F}, X)_{0\,i} < 0$ for all $i = 1, \dots, s$. Then there exist a subflag $\mathcal{F}' \subseteq \mathcal{F}$ and a subsequence $\mathbf{r}' \subseteq \mathbf{r}$, both of length $t \leqslant s$, such that*

$$S(\mathcal{F}, X, \mathbf{r}) \leqslant S(\mathcal{F}', X, \mathbf{r}'), \tag{109}$$

$$\Delta(\mathcal{F}', X)_{i-1\,i} \leqslant 0, \; i = 1, \dots, t. \tag{110}$$

In particular, $S(\mathcal{F}, X, \mathbf{r}) < 0$.

Proof. The proof is by induction on s. For $s = 1$,

$$S(\mathcal{F}, X, \mathbf{r}) = r_1 \Delta(\mathcal{F}, X)_{0\,1} < 0. \tag{111}$$

Let now $s > 1$. If for all $i = 1, \dots, s$, $\Delta(\mathcal{F}, X)_{i-1\,i} \leqslant 0$, then we are done since $\Delta(\mathcal{F}, X)_{0\,1} < 0$. Hence, we may assume that there is $i \leqslant s$ such that

$$\Delta(\mathcal{F}, X)_{j-1\,j} \leqslant 0,\ 1 \leqslant j < i, \ \text{ and } \ \Delta(\mathcal{F}, X)_{i-1\,i} > 0. \tag{112}$$

Set $\mathcal{F}'$ to be $\mathcal{F}$ without V_i and $\mathbf{r}'$ to be $\mathbf{r}$ without r_i, then

$$\begin{aligned} S(\mathcal{F}, X, \mathbf{r}) &= \sum_{j \neq i-1, i} r_j \Delta(\mathcal{F}, X)_{j-1\,j} + r_{i-1}\Delta(\mathcal{F}, X)_{i-2\,i-1} + r_i \Delta(\mathcal{F}, X)_{i-1\,i} \\ \leqslant \sum_{j \neq i-1, i} r_j \Delta(\mathcal{F}, X)_{j-1\,j} &+ r_{i-1}(\Delta(\mathcal{F}, X)_{i-2\,i-1} + \Delta(\mathcal{F}, X)_{i-1\,i}) = S(\mathcal{F}', X, \mathbf{r}'), \end{aligned} \tag{113}$$

where in the last equality we use (108). Since the length of $\mathcal{F}'$ is less than that of $\mathcal{F}$ and $\Delta(\mathcal{F}', X)_{0\,j}$ is either $\Delta(\mathcal{F}, X)_{0\,j-1}$ or $\Delta(\mathcal{F}, X)_{0\,j}$, thus, being negative, the result follows by induction. □

Let V and U be real vector spaces. For any $\mathbf{V} \in V \otimes U$, denote the subspace

$$\mathbf{V}U^* = \{\mathbf{V}\boldsymbol{\xi} \mid \boldsymbol{\xi} \in U^*\} \subseteq V, \tag{114}$$

where $\mathbf{V}\boldsymbol{\xi}$ is the convolution along U.

Lemma 13. *Let V and U be vector spaces, and X be a family of i.i.d. continuously distributed random vectors in $V \otimes U$, then*

$$\dim \sum_{\mathbf{X} \in X} \mathbf{X}U^* = \min(|X| \dim U, \dim V), \quad a.s. \tag{115}$$

Proof. Choose bases in V and U, then the elements of X read as matrices and the space $\sum_{\mathbf{X} \in X} \mathbf{X}U^*$ is spanned by the columns of all $\mathbf{X} \in X$. Since the elements of X are i.i.d. and continuously distributed, the matrix consisting of all columns of all $\mathbf{X}$-s is of maximal rank. Since it contains $|X| \dim U$ columns and $\dim V$ rows, the result follows. □

Corollary 3. *Let V and U be vector spaces, $K \subsetneq V$ be a proper subspace, and $X \subset V \otimes U$ be a family of i.i.d. continuously distributed vectors, then*

$$|X \cap K \otimes U| \leqslant \frac{\dim K}{\dim U}, \quad a.s. \tag{116}$$

Proof. Let $Y = X \cap K \otimes U$, then $E = \sum_{\mathbf{Y} \in Y} \mathbf{Y} U^* \subseteq K$ and Lemma 13 yields

$$|Y| \dim U = \dim E \leqslant \dim K. \tag{117}$$

□

Similarly to the Gaussian case, below we change the parametrization

$$f_{\mathcal{E}}(\mathbf{P}, \mathbf{Q}; X) = \widetilde{f}_{\mathcal{E}}(\mathbf{P}^{-1}, \mathbf{Q}^{-1}; X), \tag{118}$$

which does not affect the existence and uniqueness results. Let the dimensions of V and U be p and q, correspondingly, and partition $f_{\mathcal{E}}(\mathbf{P}, \mathbf{Q}; X)$ as

$$f_{\mathcal{E}}(\mathbf{P}, \mathbf{Q}; X) = -\frac{1}{p} \ln |\mathbf{P}| - \frac{1}{q} \ln |\mathbf{Q}| + \frac{1}{n} \sum_{i=1}^{n} \ln \left(\operatorname{Tr} \left(\mathbf{P} \mathbf{X}_i \mathbf{Q} \mathbf{X}_i^\top \right) \right) = f_{\mathbf{P}} + f_{\mathbf{Q}} + f_X. \tag{119}$$

Consider $f_{\mathcal{E}}(\mathbf{P}, \mathbf{Q}; X)$ over $\mathcal{M}_{\mathcal{E}}$ defined in (34), which in our new notations means that

$$\operatorname{Tr} \left(\mathbf{P}^{-1} \right) = \operatorname{Tr} \left(\mathbf{Q}^{-1} \right) = 1. \tag{120}$$

Lemma 14. *Let V and U be p and q dimensional real linear spaces, respectively, then if*

$$|X| > \frac{\max(p, q)}{\min(p, q)}, \tag{121}$$

$$f_{\mathcal{E}} \to +\infty \ \ as \ \ \mathcal{M}_{\mathcal{E}} \ni (\mathbf{P}, \mathbf{Q}) \to \partial \mathcal{M}_{\mathcal{E}}, \ \ a.s. \tag{122}$$

Proof. Assume on the contrary, that there exists a sequence $\mathbf{T} = (\mathbf{P}, \mathbf{Q}) \subset \mathcal{M}_{\mathcal{E}}$ (we omit indices for brevity) tending to $\partial \mathcal{M}_{\mathcal{E}}$ and such that $f_{\mathcal{E}}(\mathbf{T})$ is bounded. Note that due to (120) at least a part of eigenvalues of $\mathbf{P}$ and $\mathbf{Q}$ tend to $+\infty$ and the others are bounded by positive constants from below.

Let $\mathbf{P} = \sum_{j=1}^{p} \lambda_j \mathbf{y}_j \mathbf{y}_j^\top$ and $\mathbf{Q} = \sum_{i=1}^{q} \mu_i \mathbf{z}_i \mathbf{z}_i^\top$ be the spectral decompositions of $\mathbf{P}$ and $\mathbf{Q}$. Passing to a subsequence, if needed, we may suppose everything to converge here. Below we do not mention explicitly the subsequence argument while it is assumed to be utilized if necessary. Denote the sets of eigenvalues of $\mathbf{P}$ and $\mathbf{Q}$ by Λ and M, respectively. Let $\Lambda = \sqcup_{i=1}^{u+1} \Lambda_i$ and ρ be a sequence such that $\ln \lambda / \ln \rho \to r_i$ whenever $\lambda \in \Lambda_i$ and $r_1 > \cdots > r_u > r_{u+1} = 0$. In particular, $\ln \lambda \asymp r_i \ln \rho$ for $\lambda \in \Lambda_i$ (if $r_i = 0$, by this expression we mean that $\ln \lambda = o(\ln \rho)$). Define K_i to be the

space generated by the limits of eigenvectors corresponding to the values in Λ_i, hence, $V = \oplus_{i=1}^{u+1} K_i$. Now set $V_i = (\oplus_{j=1}^{i} K_j)^{\perp}$ for all $i = 0, \dots, u$. Define a flag $\mathcal{F}$ of length u as

$$\mathcal{F} = \{V \otimes U = V_0 \otimes U \supsetneq \dots \supsetneq V_u \otimes U\} \tag{123}$$

and $\mathbf{r} = \{r_1, \dots, r_u\}$.

In a similar way, let $\mathrm{M} = \sqcup_{j=1}^{v+1} \mathrm{M}_j$ and ν be such a sequence that $\ln \mu / \ln \nu \to t_j$ whenever $\mu \in \mathrm{M}_j$ and $t_1 > \dots > t_v > t_{v+1} = 0$. In particular, $\ln \mu \asymp t_i \ln \nu$ for $\mu \in \mathrm{M}_i$. The space generated by the limits of eigenvectors corresponding to M_j will be denoted by L_j, hence, $U = \oplus_{j=1}^{v+1} L_j$. Now we set $U_i = (\oplus_{j=1}^{i} L_j)^{\perp}$ for all $i = 0, \dots, v$. Define a flag $\mathcal{G}$ of length v as

$$\mathcal{G} = \{V \otimes U = V \otimes U_0 \supsetneq \dots \supsetneq V \otimes U_v\} \tag{124}$$

and $\mathbf{t} = \{t_1, \dots, t_v\}$.

Let $E_{ij} = V_{i-1} \otimes U_{j-1} \subseteq V \otimes U$, then for any $\lambda \in \Lambda_i$, $\mu \in \mathrm{M}_j$ and $\mathbf{X} \in V_{i-1} \otimes U_{j-1}$, the limit of

$$\frac{1}{\lambda}\mathbf{P}\mathbf{X}\frac{1}{\mu}\mathbf{Q} \tag{125}$$

exists and will be denoted by $\mathbf{R}_{ij}(\mathbf{X})$. By the definition, $\mathbf{R}_{ij}$ is a composition of the orthogonal projection onto E_{ij} and a positive operator on the image. Let

$$X_{ij} = X \cap V_{i-1} \otimes U_{j-1} \setminus X \cap (V_i \otimes U_{j-1} + V_{i-1} \otimes U_j)\,, \tag{126}$$

then $X = \sqcup_{i=1}^{u+1} \sqcup_{j=1}^{v+1} X_{ij}$.

We now proceed to computing the leading asymptotic terms of the summands in (119),

$$f_{\mathbf{P}} \asymp -\frac{1}{p}\sum_{i=1}^{u+1}\sum_{\lambda\in\Lambda_i} r_i \ln \rho = -\frac{1}{p}\sum_{i=1}^{u} r_i |\Lambda_i| \ln \rho = -\sum_{i=1}^{u} r_i \frac{\dim V_{i-1} - \dim V_i}{p} \ln \rho. \tag{127}$$

Similarly,

$$f_{\mathbf{Q}} \asymp -\sum_{j=1}^{v} t_j \frac{\dim U_{j-1} - \dim U_j}{q} \ln \nu. \tag{128}$$

Let $\mathbf{X} \in X_{ij}$, then for any $\lambda \in \Lambda_i$ and $\mu \in \mathrm{M}_j$, we have

$$\begin{aligned} \ln(\mathbf{P}\mathbf{X}\mathbf{Q}, \mathbf{X}) &\asymp \ln \lambda + \ln \mu + \ln(\lambda^{-1}\mathbf{P}\mathbf{X}\mu^{-1}\mathbf{Q}, \mathbf{X}) \\ &\asymp r_i \ln \rho + t_j \ln \nu + \ln(\mathbf{R}_{ij}(\mathbf{X}), \mathbf{X}) \asymp r_i \ln \rho + t_j \ln \nu. \end{aligned} \tag{129}$$

Taking this into account, we compute

$$f_X \asymp \frac{1}{|X|}\sum_{i=1}^{u+1}\sum_{j=1}^{r+1}\sum_{x\in X_{ij}}(r_i\ln\rho + t_i\ln\nu) = \frac{1}{|X|}\sum_{i=1}^{u+1}\sum_{j=1}^{r+1}|X_{ij}|(r_i\ln\rho + t_i\ln\nu). \tag{130}$$

We are interested in the asymptotic of the sum $f_{\mathbf{P}} + f_{\mathbf{Q}} + f_X$, whose leading term, when non-zero, can be written as

$$f_{\mathcal{E}} = f_{\mathbf{P}} + f_{\mathbf{Q}} + f_X \asymp A\ln\rho + B\ln\nu, \tag{131}$$

where

$$\begin{aligned} A &= -\sum_{i=1}^{u} r_i\frac{\dim V_{i-1} - \dim V_i}{p} + \frac{1}{|X|}\sum_{i=1}^{u+1}\sum_{j=1}^{r+1}|X_{ij}|r_i \\ &= -\sum_{i=1}^{u} r_i\left(\frac{\dim V_{i-1} - \dim V_i}{p} - \frac{\sum_{j=1}^{r+1}|X_{ij}|}{|X|}\right) \\ &= -\sum_{i=1}^{u} r_i\left(\frac{\dim V_{i-1}\otimes U - \dim V_i\otimes U}{pq} - \frac{|X\cap V_{i-1}\otimes U| - |X\cap V_i\otimes U|}{|X|}\right) \\ &= -S(\mathcal{F}, X, \mathbf{r}). \end{aligned} \tag{132}$$

Similar derivation yields $B = S(\mathcal{G}, X, \mathbf{t})$. Thus,

$$f_{\mathcal{E}} \asymp -S(\mathcal{F}, X, \mathbf{r})\ln\rho - S(\mathcal{G}, X, \mathbf{t})\ln\nu, \tag{133}$$

where the right-hand side is non-zero since at least one pair of eigenvalues tend to $+\infty$ due to the trace constraint (120). In addition, this implies that $\ln\rho$ and $\ln\nu$ tend to $+\infty$, and it remains to show the coefficients $S(\cdot)$ are both negative, thus guaranteeing that $f_{\mathcal{E}} \to +\infty$. By Lemma 12, it is enough to check that $\Delta(\mathcal{F}, X)_{0i} < 0$ for $i = 1, \ldots, u$ and $\Delta(\mathcal{G}, X)_{0j} < 0$ for $j = 1\ldots, v$. We have

$$\begin{aligned} \Delta(\mathcal{F}, X)_{0i} &= \frac{\dim V\otimes U - \dim V_i\otimes U}{pq} - \frac{|X| - |X\cap V_i\otimes U|}{|X|} \\ &= \frac{|X\cap V_i\otimes U|}{|X|} - \frac{\dim V_i}{p}, \end{aligned} \tag{134}$$

and we need to show

$$\frac{|X\cap V_i\otimes U|}{|X|} < \frac{\dim V_i}{p}. \tag{135}$$

By Corollary 3, we get

$$\frac{|X \cap V_i \otimes U|}{|X|} \leqslant \frac{\dim V_i}{q|X|} < \frac{\dim V_i}{\dim V}, \tag{136}$$

where the last inequality holds because $|X| > \max[p,q]/\min[p,q]$. After a similar calculation for B, we see that both A and B are negative and $f_{\mathcal{E}}(\mathbf{T};X) \to +\infty$ as $\mathbf{T} \to \partial\mathcal{M}_{\mathcal{E}}$. This contradicts the boundedness assumption on $f_{\mathcal{E}}(\mathbf{T};X)$ and completes the proof. □

Proof of Theorem 2. Since the KP constraint and the target function are convex in the Riemannian metric introduced above, the only thing we need to show is that $f_{\mathcal{E}} \to +\infty$ when we approach $\partial\mathcal{M}_{\mathcal{E}}$. Lemma 14 proves that this is a.s. true when

$$|X| > \frac{\max(p,q)}{\min(p,q)}, \tag{137}$$

and we are done. □

References

Y. I. Abramovich, N. K. Spencer, and M. D. Turley. Time-varying autoregressive (TVAR) models for multiple radar observations. *IEEE Transactions on Signal Processing*, 55(4):1298–1311, 2007.

R. Adamczak, A. Litvak, A. Pajor, and N. Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of American Mathematic Society*, 23(2): 535–561, 2010.

G. I. Allen and R. Tibshirani. Transposable regularized covariance models with an application to missing data imputation. *Annals of Applied Statistics*, 4(2):764–790, 2010.

A. Asif and J. M. F. Moura. Block matrices with L-block-banded inverse: inversion algorithms. *IEEE Transactions on Signal Processing*, 53(2):630–642, 2005.

J. Bai and S. Shi. Estimating high dimensional covariance matrices and its applications. *Annals of Economics and Finance*, 12(2):199–215, 2011.

F. Bandiera, O. Besson, and G. Ricci. Knowledge-aided covariance matrix estimation and adaptive detection in compound-Gaussian noise. *IEEE Transactions on Signal Processing*, 58(10):5391–5396, 2010.

O. Banerjee, L. El Ghaoui, and A. d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516, 2008.

P. J. Bickel and E. Levina. Regularized estimation of large covariance matrices. *Annals of Statistics*, 36(1):199–227, 2008.

E. V. Bonilla, K. M. Chai, and C. Williams. Multi-task Gaussian process prediction. *Advances in neural information processing systems*, pages 153–160, 2007.

Y. Chen, A. Wiesel, and A. O. Hero. Robust shrinkage estimation of high-dimensional covariance matrices. *IEEE Transactions on Signal Processing*, 59(9):4097–4107, 2011.

N. Cressie. Statistics for spatial data. *John Wiley and Sons*, 2015.

J. Dahmen, D. Keysers, M. Pitz, and H. Ney. Structured covariance matrices for statistical image object recognition. *Mustererkennung*, pages 99–106, 2000.

A. P. Dawid. Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika*, 68(1):265–274, 1981.

G. Derado, F. D. Bowman, and C. D. Kilts. Modeling the spatial and temporal dependence in fMRI data. *Biometrics*, 66(3):949–957, 2010.

P. Dutilleul. The MLE algorithm for the matrix normal distribution. *Journal of Statistical Computation and Simulation*, 64(2):105–123, 1999.

J. Fan, Y. Fan, and J. Lv. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, 147(1):186–197, 2008.

G. Frahm. Generalized elliptical distributions: theory and applications. *Universität zu Köln*, 2004.

G. Frahm and U. Jaekel. Tyler's M-estimator, random matrix theory, and generalized elliptical distributions with applications to finance. *Discussion papers in statistics and econometrics*, 2007.

J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.

J. H. Friedman. Regularized discriminant analysis. *Journal of the American statistical association*, 84(405):165–175, 1989.

D. R. Fuhrmann. Application of Toeplitz covariance estimation to adaptive beamforming and detection. *IEEE Transactions on Signal Processing*, 39 (10):2194–2198, 1991.

M. Greco and F. Gini. Cramér-Rao lower bounds on covariance matrix estimation for complex elliptically symmetric distributions. *IEEE Transactions on Signal Processing*, 61(24):6401–6409, 2013.

A. K. Gupta and D. K. Nagar. Matrix variate distributions. *Monographs and Surveys in Pure and Applied Mathematics, CRC Press*, 104, 1999.

A. Hero and B. Rajaratnam. Hub discovery in partial correlation graphs. *IEEE Transactions on Information Theory*, 58(9):6064–6078, 2012.

P. J. Huber. Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 35(1):73–101, 1964.

I. M. Johnstone and A. Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486), 2009.

A. Kavcic and J. M. F. Moura. Matrices with banded inverses: Inversion algorithms and factorization of Gauss-Markov processes. *IEEE Transactions on Information Theory*, 46(4):1495–1509, 2000.

E. J. Kelly. An adaptive detection algorithm. *IEEE Transactions on Aerospace and Electronic Systems*, (2):115–127, 1986.

S. L. Lauritzen. Graphical models. *Oxford University Press*, 1996.

O. Ledoit and M. Wolf. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10(5):603–621, 2003.

K. Lounici et al. High-dimensional covariance matrix estimation with missing observations. *Bernoulli*, 20(3):1029–1058, 2014.

N. Lu and D. Zimmerman. On likelihood-based inference for a separable covariance matrix. *Statistics and Actuarial Science Dept., Univ. of Iowa, Iowa City, IA, Tech. Rep*, 337, 2004.

N. Lu and D. L. Zimmerman. The likelihood ratio test for a separable covariance matrix. *Statistics and Probability Letters*, 73(4):449–457, 2005.

R. A. Maronna. Robust M-estimators of multivariate location and scatter. *Annals of Statistics*, 14(2):51–67, 1976.

E. Ollila and V. Koivunen. Robust antenna array processing using M-estimators of pseudo-covariance. *IEEE Proceedings on Personal, Indoor and Mobile Radio Communications*, 3:2659–2663, 2003.

E. Ollila, D. Tyler, V. Koivunen, and H. Poor. Complex elliptically symmetric distributions: survey, new results and applications. *IEEE Transactions on Signal Processing*, 60(11):5597–5625, 2012.

F. Pascal, Y. Chitour, J. P. Ovarlez, P. Forster, and P. Larzabal. Covariance structure maximum-likelihood estimates in compound Gaussian noise: Existence and algorithm analysis. *IEEE Transactions on Signal Processing*, 56(1):34–48, 2008.

T. Rapcsák. Geodesic convexity in nonlinear optimization. *Journal of Optimization Theory and Applications*, 69(1):169–183, 1991.

P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu. High-dimensional covariance estimation by minimizing l_1-penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.

W. J. Roberts and Y. Ephraim. Hidden Markov modeling of speech using Toeplitz covariance matrices. *Speech Communication*, 31(1):1–14, 2000.

B. Roś, F. Bijma, J. C. de Munck, and M. C. M. de Gunst. Existence and uniqueness of the maximum likelihood estimator for models with a kronecker product covariance structure. *Journal of Multivariate Analysis*, 143:345–361, 2016.

A. J. Rothman, P. J. Bickel, E. Levina, and J. Zhu. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.

J. Schäfer and K. Strimmer. A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Statistical applications in genetics and molecular biology*, 4(1), 2005.

P. Shah and V. Chandrasekaran. Group symmetry and covariance regularization. *Electronic Journal of Statistics*, 6:1600–1640, 2012.

D. L. Snyder, J. A. O'Sullivan, and M. I. Miller. The use of maximum likelihood estimation for forming images of diffuse radar targets from delay-Doppler data. *IEEE Transactions on Information Theory*, 35(3):536–548, 1989.

I. Soloveychik and A. Wiesel. Tyler's covariance matrix estimator in elliptical models with convex structure. *IEEE Transactions on Signal Processing*, 62(20):5251–5259, 2014.

I. Soloveychik and A. Wiesel. Performance analysis of Tyler's covariance estimator. *IEEE Transactions on Signal Processing*, 63(2):418–426, 2015.

I. Soloveychik and A. Wiesel. Group symmetric robust covariance estimation. *IEEE Transactions on Signal Processing*, 64(1):244–257, 2016.

M. S. Srivastava, T. von Rosen, and D. von Rosen. Models with a Kronecker product covariance structure: estimation and testing. *Mathematical Methods of Statistics*, 17(4):357–370, 2008.

Y. Sun, P. Babu, and D. P. Palomar. Robust estimation of structured covariance matrix for heavy-tailed elliptical distributions. *arXiv preprint arXiv:1506.05215*, 2015.

T. Tsiligkaridis, A. O. Hero, and S. Zhou. On convergence of Kronecker graphical lasso algorithms. *IEEE Transactions on Signal Processing*, 61 (7):1743–1755, 2013.

D. E. Tyler. A distribution-free M-estimator of multivariate scatter. *Annals of Statistics*, 15(1):234–251, 1987.

K. Werner and M. Jansson. Estimation of Kronecker structured channel covariances using training data. *Proceedings of European Signal Processing Conference*, 2007.

K. Werner, M. Jansson, and P. Stoica. On estimation of covariance matrices with Kronecker product structure. *IEEE Transactions on Signal Processing*, 56(2):478–491, 2008.

A. Wiesel. Regularized covariance estimation in scaled Gaussian models. *IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pages 309–312, 2011.

A. Wiesel. Geodesic convexity and covariance estimation. *IEEE Transactions on Signal Processing*, 60(12):6182–6189, 2012a.

A. Wiesel. Unified framework to regularized covariance estimation in scaled Gaussian models. *IEEE Transactions on Signal Processing*, 60(1):29–38, 2012b.

A. Wiesel and T. Zhang. Structured robust covariance estimation. *Foundations and Trends in Signal Processing*, 8(3):127–216, 2015.

A. Wiesel, O. Bibi, and A. Globerson. Time varying autoregressive moving average models for covariance estimation. *IEEE Transactions on Signal Processing*, 61(11):2791–2801, 2013.

J. Xie and P. M. Bentler. Covariance structure models for gene expression microarray data. *Structural Equation Modeling*, 10(4):566–582, 2003.

J. Yin and H. Li. Model selection and estimation in the matrix normal graphical model. *Journal of multivariate analysis*, 107:119–140, 2012.

K. Yu, J. Lafferty, S. Zhu, and Y. Gong. Large-scale collaborative prediction using a nonparametric random effects model. *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1185–1192, 2009.

M. Yuan and Y. Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.

Y. Zhang and J. G. Schneider. Learning multiple tasks with a sparse matrix-normal penalty. *Advances in Neural Information Processing Systems*, pages 2550–2558, 2010.